



US011348122B1

(12) **United States Patent**
Kramme et al.

(10) **Patent No.:** **US 11,348,122 B1**
(45) **Date of Patent:** ***May 31, 2022**

(54) **IDENTIFYING FRAUDULENT ONLINE APPLICATIONS**

(71) Applicant: **State Farm Mutual Automobile Insurance Company**, Bloomington, IL (US)

(72) Inventors: **Timothy Kramme**, Parker, TX (US); **Elizabeth A. Flowers**, Bloomington, IL (US); **Reena Batra**, Alpharetta, GA (US); **Miriam Valero**, Bloomington, IL (US); **Puneit Dua**, Bloomington, IL (US); **Shanna L. Phillips**, Bloomington, IL (US); **Russell Ruestman**, Minonk, IL (US); **Bradley A. Craig**, Normal, IL (US)

(73) Assignee: **State Farm Mutual Automobile Insurance Company**, Bloomington, IL (US)

(*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 0 days.

This patent is subject to a terminal disclaimer.

(21) Appl. No.: **16/913,814**

(22) Filed: **Jun. 26, 2020**

Related U.S. Application Data

(63) Continuation of application No. 15/465,874, filed on Mar. 22, 2017, now Pat. No. 10,825,028.

(Continued)

(51) **Int. Cl.**
G06Q 30/00 (2012.01)
G06Q 20/40 (2012.01)

(Continued)

(52) **U.S. Cl.**
CPC **G06Q 30/0185** (2013.01); **G06Q 20/4016** (2013.01); **G06N 20/00** (2019.01); **G06Q 30/0248** (2013.01)

(58) **Field of Classification Search**

CPC G06Q 20/4016; G06Q 30/0185; G06Q 30/0248

See application file for complete search history.

(56) **References Cited**

U.S. PATENT DOCUMENTS

5,467,341 A 11/1995 Matsukane et al.
5,708,422 A 1/1998 Blonder et al.

(Continued)

FOREIGN PATENT DOCUMENTS

CA 2741408 A1 * 4/2010 G06Q 20/102
WO WO-2006021088 A1 * 3/2006 G06F 21/604

OTHER PUBLICATIONS

R. E. Neira, "Identity Theft: Inside the Mind of a Cybercriminal." Order No. 10109629, Utica College, Ann Arbor, May 2016. (Year: 2016).*

(Continued)

Primary Examiner — Gerardo Araque, Jr.

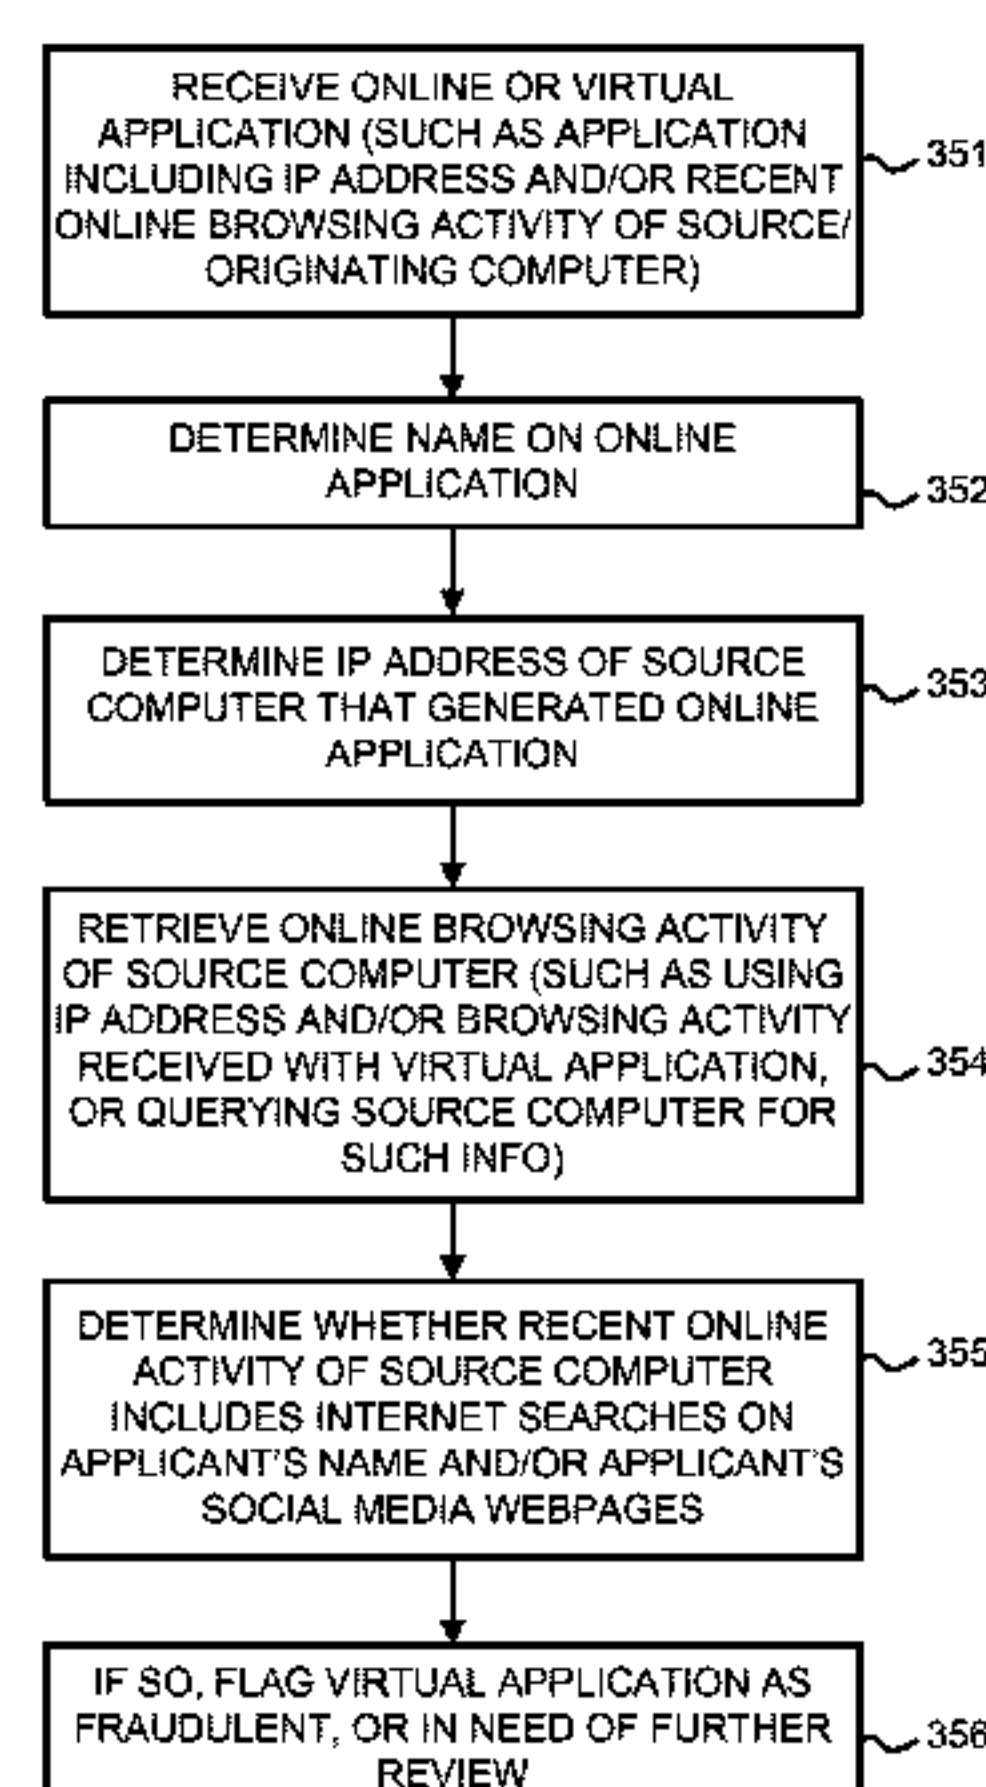
(74) *Attorney, Agent, or Firm* — Lee & Hayes, P.C.

(57) **ABSTRACT**

A method of using browsing activity to identify fraudulent online or virtual applications includes receiving a virtual application over one or more radio frequency links, determining an applicant name on the virtual application, determining an IP address of a source computer from which the virtual application originated, determining an online browsing or search history associated with the IP address, determining whether the online browsing or search history indicates recent Internet searches for the applicant name, and, in response to determining that the online browsing or search history does indicate recent Internet searches for the applicant name, flagging the virtual application as fraudulent and generating an electronic alert indicating that the virtual application is fraudulent to facilitate identifying fraudulent virtual applications for goods or services.

20 Claims, 22 Drawing Sheets

350



Related U.S. Application Data						
(60)	Provisional application No. 62/365,699, filed on Jul. 22, 2016, provisional application No. 62/331,530, filed on May 4, 2016, provisional application No. 62/318,423, filed on Apr. 5, 2016, provisional application No. 62/313,196, filed on Mar. 25, 2016.		9,485,265	B1	11/2016	Saperstein
			9,519,903	B2	12/2016	Kannan et al.
			9,569,767	B1	2/2017	Lewis et al.
			9,607,318	B1	3/2017	Gerchikov et al.
			9,691,066	B2	6/2017	McGuinness et al.
			9,721,253	B2	8/2017	Gideoni, et al.
			9,779,403	B2	10/2017	Ranganath et al.
			9,786,015	B1	10/2017	Roumeliotis
			9,818,105	B2	11/2017	Jung et al.
			9,858,575	B2	1/2018	Meredith et al.
(51)	Int. Cl. <i>G06N 20/00</i> (2019.01) <i>G06Q 30/02</i> (2012.01)	9,883,040	B2	1/2018	Strong et al.	
		9,916,606	B2	3/2018	Stroh	
		9,934,498	B2	4/2018	Jung et al.	
		9,947,055	B1	4/2018	Roumeliotis	
(56)	References Cited U.S. PATENT DOCUMENTS	9,953,326	B2	4/2018	Plymouth et al.	
		10,013,646	B2	7/2018	Ching	
		10,037,532	B2	7/2018	Birukov et al.	
		10,043,071	B1	8/2018	Wu	
		10,169,761	B1 *	1/2019	Burger G06Q 20/4016	
		10,373,140	B1	8/2019	Chang et al.	
		10,373,160	B2	8/2019	Ranganathan	
		10,565,585	B2	2/2020	Sharan et al.	
		11,170,375	B1 *	11/2021	Kramme G06N 5/046	
		2002/0099649	A1	7/2002	Lee et al.	
			2002/0133721	A1	9/2002	Adjaoute
			2002/0147694	A1	10/2002	Dempsey et al.
			2002/0169717	A1	11/2002	Challener
			2003/0033228	A1	2/2003	Bosworth-Davies et al.
			2003/0144952	A1	7/2003	Brown et al.
			2003/0172036	A1	9/2003	Feigenbaum
			2003/0182194	A1	9/2003	Choey et al.
			2004/0085195	A1	5/2004	McKibbin
			2004/0254868	A1	12/2004	Kirkland et al.
			2005/0097046	A1	5/2005	Singfield
			2005/0097051	A1	5/2005	Madill et al.
			2005/0137982	A1	6/2005	Michelassi et al.
			2005/0154676	A1	7/2005	Ronning et al.
			2005/0182712	A1	8/2005	Angell
			2005/0188349	A1	8/2005	Bent et al.
			2006/0041464	A1	2/2006	Powers et al.
			2006/0041506	A1	2/2006	Mason et al.
			2006/0041508	A1	2/2006	Pham et al.
			2006/0065717	A1	3/2006	Hurwitz et al.
			2006/0282660	A1	12/2006	Varghese et al.
			2007/0094137	A1	4/2007	Phillips et al.
			2007/0100773	A1	5/2007	Wallach
			2007/0106582	A1	5/2007	Baker et al.
			2007/0179849	A1	8/2007	Jain
			2007/0192240	A1	8/2007	Crooks
			2007/0226095	A1	9/2007	Petriuc
			2007/0244782	A1	10/2007	Chimento
			2007/0275399	A1	11/2007	Lathrop et al.
			2007/0288641	A1	12/2007	Lee et al.
			2008/0059351	A1	3/2008	Richey et al.
			2008/0106726	A1	5/2008	Park
			2008/0288299	A1	11/2008	Schultz
			2008/0290154	A1	11/2008	Barnhardt et al.
			2009/0018934	A1	1/2009	Peng et al.
			2009/0018940	A1	1/2009	Wang et al.
			2009/0044279	A1	2/2009	Crawford et al.
			2009/0083080	A1	3/2009	Brooks
			2009/0099884	A1	4/2009	Hoefelmeyer et al.
			2009/0106134	A1	4/2009	Royyuru
			2009/0192855	A1	7/2009	Subramanian et al.
			2009/0204254	A1	8/2009	Weber
			2009/0323972	A1	12/2009	Kohno et al.
			2010/0023455	A1	1/2010	Dispensa et al.
			2010/0023467	A1	1/2010	Iwakura et al.
			2010/0094767	A1	4/2010	Miltonberger
			2010/0114774	A1	5/2010	Linaman et al.
			2010/0274691	A1	10/2010	Hammad et al.
			2010/0293090	A1	11/2010	Domenikos et al.
			2010/0320266	A1	12/2010	White
			2011/0055074	A1	3/2011	Chen et al.
			2011/0066493	A1	3/2011	Faith et al.
			2011/0184845	A1	7/2011	Bishop
			2011/0208601	A1	8/2011	Ferguson et al.
			2011/0225064	A1	9/2011	Fou
			2011/0238510	A1	9/2011	Rowen et al.

(56)

References Cited**U.S. PATENT DOCUMENTS**

2011/0238566 A1 9/2011 Santos
 2011/0238575 A1 9/2011 Nightengale et al.
 2011/0258118 A1 10/2011 Ciurea
 2011/0264612 A1 10/2011 Ryman-Tubb
 2012/0036037 A1 2/2012 Xiao et al.
 2012/0047072 A1 2/2012 Larkin
 2012/0070062 A1 3/2012 Houle et al.
 2012/0109821 A1 5/2012 Barbour et al.
 2012/0158566 A1 6/2012 Fok et al.
 2012/0173570 A1 7/2012 Golden
 2012/0209773 A1 8/2012 Ranganathan
 2012/0214442 A1 8/2012 Crawford et al.
 2012/0216260 A1 8/2012 Crawford et al.
 2012/0226613 A1 9/2012 Adjaoute
 2012/0245967 A1 9/2012 Dispensa et al.
 2012/0278249 A1* 11/2012 Duggal G06Q 40/08
 705/325
 2013/0013491 A1 1/2013 Selway et al.
 2013/0018781 A1 1/2013 Prada Peyser et al.
 2013/0024339 A1 1/2013 Choudhuri et al.
 2013/0024358 A1 1/2013 Choudhuri et al.
 2013/0024373 A1 1/2013 Choudhuri et al.
 2013/0046692 A1 2/2013 Grigg et al.
 2013/0085942 A1 4/2013 Shirol
 2013/0197998 A1* 8/2013 Buhrmann G06Q 20/4016
 705/14.53
 2013/0218758 A1 8/2013 Koenigsbrueck et al.
 2013/0262311 A1 10/2013 Buhrmann et al.
 2013/0290119 A1 10/2013 Howe et al.
 2013/0297501 A1 11/2013 Monk et al.
 2014/0012701 A1 1/2014 Wall et al.
 2014/0012738 A1 1/2014 Woo
 2014/0052621 A1 2/2014 Love
 2014/0058962 A1 2/2014 Davoodi et al.
 2014/0067656 A1 3/2014 Cohen Ganor et al.
 2014/0074762 A1 3/2014 Campbell
 2014/0101050 A1 4/2014 Clarke et al.
 2014/0122325 A1 5/2014 Zoldi et al.
 2014/0200929 A1* 7/2014 Fitzgerald H04W 8/22
 705/4
 2014/0201126 A1 7/2014 Zadeh et al.
 2014/0207637 A1 7/2014 Groarke
 2014/0207674 A1 7/2014 Schroeder et al.
 2014/0244317 A1 8/2014 Roberts et al.
 2014/0244503 A1 8/2014 Sadlier
 2014/0250011 A1 9/2014 Weber
 2014/0266669 A1 9/2014 Fadell et al.
 2014/0279312 A1 9/2014 Mason et al.
 2014/0279494 A1 9/2014 Wiesman et al.
 2014/0279503 A1 9/2014 Bertanzetti et al.
 2014/0282856 A1 9/2014 Duke et al.
 2014/0310160 A1 10/2014 Kumar et al.
 2014/0310176 A1 10/2014 Saunders et al.
 2014/0324677 A1 10/2014 Walraven et al.
 2014/0337217 A1 11/2014 Howe et al.
 2014/0337243 A1 11/2014 Dutt et al.
 2014/0358592 A1 12/2014 Wedig et al.
 2014/0373114 A1 12/2014 Franca-Neto et al.
 2015/0012436 A1 1/2015 Poole et al.
 2015/0032604 A1 1/2015 Kearns et al.
 2015/0032621 A1 1/2015 Kar et al.
 2015/0039504 A1 2/2015 Ebbert
 2015/0039513 A1 2/2015 Adjaoute
 2015/0046216 A1 2/2015 Adjaoute
 2015/0046220 A1 2/2015 Gerard et al.
 2015/0058119 A1 2/2015 Atli et al.
 2015/0081349 A1 3/2015 Johndrow et al.
 2015/0089615 A1 3/2015 Krawczyk et al.
 2015/0106260 A1 4/2015 Andrews et al.
 2015/0106268 A1 4/2015 Carroll et al.
 2015/0120502 A1 4/2015 Jung et al.
 2015/0142595 A1 5/2015 Acuna-Rohter
 2015/0148061 A1 5/2015 Koukoumidis et al.
 2015/0161611 A1 6/2015 Duke et al.
 2015/0178715 A1 6/2015 Buhrmann et al.

2015/0178733 A1 6/2015 Kozloski et al.
 2015/0186888 A1 7/2015 Katz et al.
 2015/0186891 A1 7/2015 Wagner et al.
 2015/0193768 A1 7/2015 Douglas et al.
 2015/0199699 A1 7/2015 Milton et al.
 2015/0199738 A1 7/2015 Jung et al.
 2015/0213246 A1 7/2015 Turgeman et al.
 2015/0227934 A1 8/2015 Chauhan
 2015/0227935 A1 8/2015 Adjaoute
 2015/0235221 A1 8/2015 Murphy, Jr. et al.
 2015/0242856 A1 8/2015 Dhurandhar et al.
 2015/0254659 A1 9/2015 Kulkarni et al.
 2015/0262184 A1 9/2015 Wang et al.
 2015/0269578 A1 9/2015 Subramanian et al.
 2015/0276237 A1 10/2015 Daniels et al.
 2015/0365388 A1* 12/2015 Little G06F 21/34
 726/7
 2015/0379430 A1 12/2015 Dirac et al.
 2015/0379517 A1 12/2015 Jin et al.
 2016/0044048 A1 2/2016 Hamidi et al.
 2016/0055568 A1 2/2016 Vidal et al.
 2016/0063501 A1 3/2016 Kalyan et al.
 2016/0063502 A1 3/2016 Adjaoute
 2016/0071017 A1 3/2016 Adjaoute
 2016/0071105 A1 3/2016 Groarke et al.
 2016/0078014 A1 3/2016 Avasarala et al.
 2016/0078443 A1 3/2016 Tomasofsky et al.
 2016/0086184 A1 3/2016 Carpenter et al.
 2016/0086185 A1 3/2016 Adjaoute
 2016/0104163 A1 4/2016 Aquino et al.
 2016/0132851 A1 5/2016 Desai et al.
 2016/0132882 A1 5/2016 Choudhuri et al.
 2016/0132886 A1 5/2016 Burke et al.
 2016/0140561 A1 5/2016 Cowan
 2016/0140562 A1 5/2016 Birukov et al.
 2016/0148127 A1 5/2016 Harkey et al.
 2016/0162895 A1 6/2016 Nuzum et al.
 2016/0171498 A1 6/2016 Wang et al.
 2016/0171499 A1 6/2016 Meredith et al.
 2016/0203490 A1 7/2016 Gupta et al.
 2016/0210631 A1 7/2016 Ramasubramanian et al.
 2016/0210857 A1 7/2016 Gao et al.
 2016/0232465 A1 8/2016 Kurtz et al.
 2016/0247143 A1 8/2016 Ghosh
 2016/0253731 A1 9/2016 Ketchel, III et al.
 2016/0283945 A1 9/2016 Gonzalez
 2016/0300214 A1 10/2016 Chaffin et al.
 2016/0342963 A1 11/2016 Zoldi et al.
 2017/0011382 A1 1/2017 Zoldi et al.
 2017/0070484 A1 3/2017 Kruse et al.
 2017/0103398 A1 4/2017 Napsky et al.
 2017/0178134 A1 6/2017 Senci et al.
 2017/0193727 A1 7/2017 Van Horn et al.
 2017/0201498 A1 7/2017 Baig et al.
 2017/0221062 A1 8/2017 Katz et al.
 2017/0236125 A1 8/2017 Guise et al.
 2017/0243220 A1 8/2017 Phillips et al.
 2017/0261852 A1 9/2017 Kato et al.
 2017/0270526 A1 9/2017 Fitzgerald
 2017/0286962 A1 10/2017 Lai et al.
 2017/0323345 A1 11/2017 Flowers et al.
 2018/0053114 A1 2/2018 Adjaoute
 2018/0060839 A1 3/2018 Murali
 2018/0158062 A1 6/2018 Kohli
 2018/0182029 A1 6/2018 Vinay
 2018/0232605 A1 8/2018 Chen et al.
 2019/0188212 A1 6/2019 Miller et al.
 2021/0065186 A1 3/2021 Kramme et al.
 2021/0158355 A1* 5/2021 Kramme G06V 30/40
 2021/0264429 A1* 8/2021 Kramme G06Q 20/102
 2021/0264458 A1* 8/2021 Kramme G06V 30/194
 2021/0374753 A1* 12/2021 Kramme G06Q 20/4016
 2021/0374764 A1 12/2021 Kramme et al.

OTHER PUBLICATIONS

Anonymous "Senate Bill aids in cyber fraud prosecution," Toledo Business Journal, vol. 28, (1), pp. 9, Jan. 2012 (Year: 2012).*

(56)

References Cited

OTHER PUBLICATIONS

A. Rege, "What's Love Got to Do with It? Exploring Online Dating Scams and Identity Fraud," *International Journal of Cyber Criminology*, vol. 3, (2), pp. 494-512, Jul.-Dec. 2009. (Year: 2009).*

O. Angelopoulou, "Analysis of Digital Evidence in Identity Theft Investigations." Order No. U580269, University of South Wales (United Kingdom), Ann Arbor, 2010. (Year: 2010).*

D. R. Cornelius, "Online Identity Theft Victimization: An Assessment of Victims and Non-Victims Level of Cyber Security Knowledge." Order No. 10253894, Colorado Technical University, Ann Arbor, 2016. (Year: 2016).*

Non Final Office Action dated Jul. 14, 2020 for U.S. Appl. No. 15/465,838 "Detecting Financial Fraud Using Spending Pattern Data for Individual Cardholders" Kramme, 56 pages.

Office Action for U.S. Appl. No. 15/466,002, dated Apr. 8, 2021, Kramme, "Reducing False Positive Fraud Alerts for Card-Present Financial Transactions", 75 pages.

Office Action for U.S. Appl. No. 15/465,821, dated May 12, 2021, Kramme, "Preempting or Resolving Fraud Disputes Relating to Introductory Offer Expirations", 36 pages.

Office Action dated Mar. 16, 2021 for U.S. Appl. No. 15/465,821 "Reducing False Positives Based Upon Customer Online Activity" Kramme, 20 pages.

Carneiro, "A Data Mining Approach to Fraud Detection in e-tail", Integrated Master in Industrial Engineering and Management, Jan. 2016, 71 pages.

Final Office Action dated Dec. 15, 2020 for U.S. Appl. No. 15/466,014, "Identifying False Positive Geolocation-Based Fraud Alerts", Kramme, 25 pages.

Non Final Office Action dated Dec. 22, 2020 for U.S. Appl. No. 15/465,856, "Identifying Potential Chargeback Scenarios Using Machine Learning", Kramme, 23 pages.

Non Final Office Action dated Dec. 22, 2020 for U.S. Appl. No. 15/465,880, "Preempting or Resolving Fraud Disputes Relating to Introductory Offer Expirations", Kramme, 58 pages.

Non Final Office Action dated Dec. 22, 2020 for U.S. Appl. No. 15/466,002, "Reducing False Positive Fraud Alerts for Card-Present Financial Transactions", Kramme, 53 pages.

Final Office Action dated Dec. 30, 2020 for U.S. Appl. No. 15/465,838, "Detecting Financial Fraud Using Spending Pattern Data for Individual Cardholders", Kramme, 60 pages.

sas.com, "Machine Learning What it is & Why it Matters", retrieved from <<https://www.sas.com/en_us/insights/analytics/machine-learning.html>>, Mar. 18, 2016, 11 pages.

Office Action for U.S. Appl. No. 15/465,868, dated Jun. 17, 2021, Kramme, "Facilitating Fraud Dispute Resolution Using Machine Learning", 29 Pages.

Office Action for U.S. Appl. No. 15/465,856, dated Jul. 6, 2021, Kramme, "Identifying Potential Chargeback Scenarios Using Machine Learning", 25 Pages.

Rowan, et al., "Evaluation of Machine Learning Tools for Distinguishing Fraud from Error", *Journal of Business & Economics Research*—Sep. 2013, vol. 11, No. 9, 2013, pp. 393-400.

Yair Finzi, "Fraud Losses and False Positives: The Numbers", Secured Touch, retrieved at <<<https://blog.securedtouch.com/fraud-losses-and-positives-numbers>>>, 2015, 9 pgs.

Yair Finzi, "Why Businesses are Struggling to Fight Mobile Fraud?", Secured Touch, retrieved at <<<https://blog.securedtouch.com/why-businesses-are-struggling-to-fight-mobile-fraud>>>, 2015, 8 pgs.

Abdelhalim, et al., "The Impact of Google Hacking on Identity and Application Fraud", 2007 IEEE Pacific Rim Conference on Communications, Computers and Signal Processing, Aug. 2007, pp. 240-244.

Aimeur, et. al. "The ultimate invasion of privacy: Identity theft", Ninth Annual International Conference on Privacy, Security, and Trust, Montreal, Quebec, 2011, pp. 24-31.

Bhatia T.P., "Understanding Credit Card Frauds", *Cards Business Review* #2003-01, Jun. 2003, 17 pages.

Bose, "Intelligent Technologies for Managing Fraud and Identity Theft", Third International Conference on Information Technology: New Generations (ITNG'06), Apr. 2006, pp. 446-451.

Cho, et. al., "Detection and Response of Identity Theft within a Company Utilizing Location Information," International Conference on Platform Technology and Service (PlatCon), Feb. 2016, 5 pages.

Council, "How to Extract Reliable Data from Invoices" retrieved on Dec. 5, 2019, at <<<https://www.parascript.com/blog/how-to-extract-reliable-date-from-invoices/>>>Parascript blog, Mar. 10, 2016, pp. 1-4.

Credit Card Accountability Responsibility and Disclosure Act of 2009, 123 Stat. 1734, Public Law 111-24, 111th Congress, May 22, 2009 (Year: 2009).

Expert System, "What is Machine Learning? A definition", Mar. 7, 2017, retrieved on Jan. 31, 2020: <<<https://expertsystem.com/machine-learning-definition/>>>, 2 pages.

Gold, "Identity Crisis?", *Engineering & Technology*, vol. 8, No. 10, Nov. 2013, pp. 32-35.

Ji, et. al., "Systems Plan for Combating Identity Theft—A Theoretical Framework," 2007 International Conference on Wireless Communications, Networking and Mobile Computing, Shanghai, Sep. 2007, pp. 6402-6405.

Kejriwal, et al., "Semi-supervised Instance Matching Using Boosted Classifiers", *The Semantic Web, Latest Advances and New Domains, ESWC 2015, Lecture Notes in Computer Science*, vol. 9088, May 2015, pp. 1-15.

Kossman, "7 merchant tips to understanding EMV fraud liability shift", *CreditCards.com*, Oct. 1, 2015, retrieved from <<<https://www.creditcards.com/credit-card-news/understanding-emv-fraud-liability-shift-1271.php>>>, 14 pages.

ProgrammableWeb website, "MasterCard Merchant Identifier API", Oct. 14, 2015, <<<https://web.archive.org/web/20151014154922/http://www.programmableweb.com/api/mastercard-merchant-identifier>>>, 3 pages.

Van der Merwe, et. al., "Phishing in the System of Systems Settings: Mobile Technology", 2005 IEEE International Conference on Systems, Man and Cybernetics, Oct. 2005, vol. 1, pp. 492-498.

Office Action for U.S. Appl. No. 15/465,868, dated Feb. 20, 2020, Kramme, "Facilitating Fraud Dispute Resolution Using Machine Learning", 39 pages.

Office Action for U.S. Appl. No. 15/465,838, dated Mar. 5, 2020, Kramme, "Detecting Financial Fraud Using Spending Pattern Data for Individual Cardholders", 52 pages.

Final Office Action dated Jun. 2, 2020 for U.S. Appl. No. 15/465,858, "Identifying Chargeback Scenarios Based Upon Non-Compliant Merchant Computer Terminals", Kramme, 13 pages.

Final Office Action dated Jun. 10, 2020 for U.S. Appl. No. 15/465,880, "Preempting or Resolving Fraud Disputes Relating to Introductory Offer Expirations", Kramme, 37 pages.

Office Action for U.S. Appl. No. 15/466,014, dated Jan. 25, 2019, Kramme, "Identifying False Positive Geolocation-Based Fraud Alerts", 32 pages.

Office Action for U.S. Appl. No. 15/466,002, dated Oct. 16, 2019, Kramme, "Reducing False Positive Fraud Alerts for Card-Present Financial Transactions", 26 pages.

Office Action for U.S. Appl. No. 15/465,871, dated Oct. 31, 2019, Kramme, "Preempting or Resolving Fraud Disputes Relating to Billing Aliases", 31 pages.

Office Action for U.S. Appl. No. 15/465,827, dated Oct. 9, 2019, Kramme, "Reducing False Positives Using Customer Data and Machine Learning", 36 pages.

Office Action for U.S. Appl. No. 15/465,868, dated Nov. 19, 2018, Kramme, "Facilitating Fraud Dispute Resolution Using Machine Learning", 29 pages.

Office Action for U.S. Appl. No. 15/465,842, dated Nov. 27, 2019, Kramme, "Automated Fraud Classification Using Machine Learning", 30 pages.

Office Action for U.S. Appl. No. 15/465,880, dated Dec. 12, 2019, Kramme, "Preempting or Resolving Fraud Disputes Relating to Introductory Offer Expirations," 39 pages.

Office Action for U.S. Appl. No. 15/465,977, dated Dec. 12, 2019, Kramme, "Document-Based Fraud Detection," 10 pages.

(56)

References Cited

OTHER PUBLICATIONS

Office Action for U.S. Appl. No. 15/465,856, dated Dec. 18, 2019, Kramme, "Identifying Potential Chargeback Scenarios Using Machine Learning," 22 pages.

Office Action for U.S. Appl. No. 15/465,863, dated Dec. 18, 2019, Kramme, "Identifying Application-Related Fraud," 17 pages.

Office Action for U.S. Appl. No. 15/465,821, dated Dec. 19, 2019, Kramme, "Reducing False Positives Based Upon Customer Online Activity," 22 pages.

Office Action for U.S. Appl. No. 15/466,002, dated Dec. 27, 2018, Kramme, "Reducing False Positive Fraud Alerts for Card-Present Financial Transactions," 35 pages.

Office Action for U.S. Appl. No. 15/466,009, dated Dec. 27, 2018, Kramme, "Reducing False Positive Fraud Alerts for Online Financial Transactions," 33 pages.

Office Action for U.S. Appl. No. 15/465,838, dated Dec. 4, 2019, Kramme, "Detecting Financial Fraud Using Spending Pattern Data for Individual Cardholders," 45 pages.

Office Action for U.S. Appl. No. 15/465,981, dated Feb. 20, 2019, Kramme, "Identifying Fraudulent Instruments and Identification," 41 pages.

Office Action for U.S. Appl. No. 15/466,002, dated Feb. 21, 2020, Kramme, "Reducing False Positive Fraud Alerts for Card-Present Financial Transactions," 29 pages.

Office Action for U.S. Appl. No. 15/465,977, dated Feb. 25, 2019, Kramme, "Document-Based Fraud Detection," 8 pages.

Office Action for U.S. Appl. No. 15/465,858, dated Feb. 3, 2020, Kramme, "Identifying Chargeback Scenarios Based Upon Non-Compliant Merchant Computer Terminals," 48 pages.

Office Action for U.S. Appl. No. 15/465,827, dated Feb. 5, 2020, Kramme, "Reducing False Positives Using Customer Data and Machine Learning," 34 pages.

Office Action for U.S. Appl. No. 15/466,009, dated Mar. 12, 2020, Kramme, "Reducing False Positive Fraud Alerts for Online Financial Transactions," 24 pages.

Office Action for U.S. Appl. No. 15/465,838, dated Mar. 20, 2019, Kramme, "Detecting Financial Fraud Using Spending Pattern Data for Individual Cardholders," 28 pages.

Office Action for U.S. Appl. No. 15/465,821, dated Mar. 5, 2019, Kramme, "Reducing False Positives Based Upon Customer Online Activity," 21 pages.

Office Action for U.S. Appl. No. 15/465,880, dated Apr. 15, 2019, Kramme, "Preempting or Resolving Fraud Disputes Relating to Introductory Offer Expirations," 25 pages.

Office Action for U.S. Appl. No. 15/465,868, dated Apr. 18, 2019, Kramme, "Facilitating Fraud Dispute Resolution Using Machine Learning," 29 pages.

Office Action for U.S. Appl. No. 15/465,871, dated Apr. 24, 2019, Kramme, "Preempting or Resolving Fraud Disputes Relating to Billing Aliases," 35 pages.

Office Action for U.S. Appl. No. 15/465,863, dated Apr. 28, 2020, Kramme, "Identifying Application-Related Fraud," 14 pages.

Office Action for U.S. Appl. No. 15/465,871, dated Apr. 3, 2020, Kramme, "Preempting or Resolving Fraud Disputes Relating to Billing Aliases," 38 pages.

Office Action for U.S. Appl. No. 15/465,863, dated Apr. 30, 2019, Kramme, "Identifying Application-Related Fraud," 19 pages.

Office Action for U.S. Appl. No. 15/466,014, dated Apr. 8, 2020, Kramme, "Identifying False Positive Geolocation-Based Fraud Alerts," 26 pages.

Office Action for U.S. Appl. No. 15/465,821, dated Apr. 9, 2020, Kramme, "Reducing False Positives Based Upon Customer Online Activity," 23 pages.

Office Action for U.S. Appl. No. 15/465,977, dated Apr. 9, 2020, Kramme, "Document-Based Fraud Detection," 10 pages.

Office Action for U.S. Appl. No. 15/465,827, dated May 1, 2019, Kramme, "Reducing False Positives Using Customer Data and Machine Learning," 32 pages.

Final Office Action dated May 21, 2020 for U.S. Appl. No. 15/465,856 "Identifying Potential Chargeback Scenarios Using Machine Learning" Kramme, 22 pages.

Final Office Action dated May 28, 2020 for U.S. Appl. No. 15/465,842 "Automated Fraud Classification Using Machine Learning" Kramme, 32 pages.

Office Action for U.S. Appl. No. 15/465,977, dated Jun. 11, 2019, Kramme, "Document-Based Fraud Detection", 10 pages.

Office Action for U.S. Appl. No. 15/466,009, dated Jun. 17, 2019, Kramme, "Reducing False Positive Fraud Alerts for Online Financial Transactions", 25 pages.

Office Action for U.S. Appl. No. 15/465,821, date Jun. 27, 2019, Kramme, "Reducing False Positives Based Upon Customer Online Activity", 24 pages.

Office Action for U.S. Appl. No. 15/466,002, dated Jun. 6, 2019, Kramme, "Reducing False Positive Fraud Alerts for Card-Present Financial Transactions", 25 pages.

Office Action for U.S. Appl. No. 15/465,838, dated Jul. 1, 2019, Kramme, "Detecting Financial Fraud Using Spending Pattern Data for Individual Cardholders", 39 pages.

Office Action for U.S. Appl. No. 15/466,014, dated Jul. 12, 2019, Kramme, "Identifying False Positive Geolocation-Based Fraud Alerts", 25 pages.

Office Action for U.S. Appl. No. 15/465,880, dated Aug. 20, 2019, Kramme, "Preempting or Resolving Fraud Disputes Relating to Introductory Offer Expirations", 27 pages.

Office Action for U.S. Appl. No. 15/465,863, dated Aug. 30, 2019, Kramme, "Identifying Application-Related Fraud", 21 pages.

Office Action for U.S. Appl. No. 15/465,981, dated Aug. 6, 2019, Kramme, "Identifying Fraudulent Instruments and Identification", 42 pages.

Rong, et al. "A Machine Learning Approach for Instance Matching Based on Similarity Metrics," International Semantic Web Conference, ISWC 2012: The Semantic Web, Nov. 2012, 16 pages.

Singh, et al., "Fraud Detection by Monitoring Customer Behavior and Activities", International Journal of Computer Applications, vol. 111, No. 11, Feb. 2015, pp. 23-32.

"Tips for Managing Rewards Credit Cards", Apr. 13, 2014, Frequent Flyer Tales blog, retrieved from Internet Archive Wayback Machine: <https://web.archive.org/web/20140413142932/https://www.rewardexpert.com/blog/tips-for-managing-rewards-credit-cards/> (Year: 2014).

Wang, et al., "A Brief Review of Machine Learning and its Application", International Conference on Information Engineering and Computer Science, Dec. 2009, 4 pages.

Office Action for U.S. Appl. No. 16/988,157, dated Aug. 18, 2021, Kramme, "Preempting or Resolving Fraud Disputes Relating to Billing Aliases", 28 Pages.

Final Office Action dated Nov. 12, 2020 for U.S. Appl. No. 15/465,981, "Identifying Fraudulent Instruments and Identification", Kramme, 33 pages.

Final Office Action dated Nov. 13, 2020 for U.S. Appl. No. 15/465,871, "Preempting or Resolving Fraud Disputes Relating to Billing Aliases", Kramme, 26 pages.

Non Final Office Action dated Nov. 17, 2020 for U.S. Appl. No. 15/465,821, "Reducing False Positives Based Upon Customer Online Activity", Kramme, 19 pages.

ProgrammableWeb, "MasterCard Merchant Identifier API", retrieved from <<<https://web.archive.org/web/20151014154922/http://www.programmableweb.com/api/mastercard-merchant-identifier>>>, Oct. 14, 2015, 3 pages.

Non Final Office Action dated Oct. 20, 2020 for U.S. Appl. No. 16/899,486, "Reducing False Positives Using Customer Data and Machine Learning", Kramme, 11 pages.

Office Action dated Oct. 29, 2020 for U.S. Appl. No. 15/465,858 "Identifying Chargeback Scenarios Based Upon Non-Compliant Merchant Computer Terminals" Kramme, 14 pages.

Office Action dated Oct. 29, 2020 for U.S. Appl. No. 15/931,560, "Identifying Chargeback Scenarios Based Upon Non-Compliant Merchant Computer Terminals", Kramme, 15 pages.

Final Office Action dated Jul. 23, 2020 for U.S. Appl. No. 15/465,868, "Facilitating Fraud Dispute Resolution Using Machine Learning", Kramme, 68 pages.

(56)

References Cited

OTHER PUBLICATIONS

Final Office Action dated Jul. 24, 2020 for U.S. Appl. No. 15/931,560, "Identifying Chargeback Scenarios Based Upon Non-Compliant Merchant Computer Terminals", Kramme, 61 pages.
 Final Office Action dated Aug. 24, 2020 for U.S. Appl. No. 15/466,009, "Reducing False Positive Fraud Alerts for Online Financial Transactions", Kramme, 27 pages.
 Non Final Office Action dated Aug. 26, 2020 for U.S. Appl. No. 15/466,014, "Identifying False Positive Geolocation-Based Fraud Alerts", Kramme, 26 pages.
 Non Final Office Action dated Aug. 27, 2020 for U.S. Appl. No. 16/540,505, "Reducing False Positives Using Customer Feedback and Machine Learning", Kramme, 33 pages.
 Non Final Office Action dated Jul. 21, 2020 for U.S. Appl. No. 15/465,981 "Identifying Fraudulent Instruments and Identification" Kramme, 33 pages.
 Cheney, "Identity Theft: Do Definitions Still Matter?", Federal Reserve Bank of Philadelphia, Payment Card Center Discussion Paper, Aug. 2005, 22 pages.
 Office Action dated Jan. 7, 2021 for U.S. Appl. No. 15/465,868, Kramme, "Facilitating Fraud Dispute Resolution Using Machine Learning", 16 pages.
 Final Office Action dated Feb. 20, 2021 for U.S. Appl. No. 15/465,842 "Automated Fraud Classification Using Machine Learning" Kramme, 23 pages.

Fawcett, et al., "Adaptive Fraud Detection", Data Mining, and Knowledge Discovery 1, retrieved at <<<http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.123.1281&rep=rep1&type=pdf>>>, 1997, pp. 291-316.
 Office Action for U.S. Appl. No. 15/465,838, dated Sep. 16, 2021, Kramme, "Detecting Financial Fraud Using Spending Pattern Data for Individual Cardholders", 45 Pages.
 Office Action for U.S. Appl. No. 17/080,476, dated Jan. 18, 2022, Kramme, "Reducing False Positives Using Customer Feedback and Machine Learning", 24 pages.
 Office Action for U.S. Appl. No. 15/465,981, dated Jan. 24, 2022, Kramme, "Preempting or Resolving Fraud Disputes Relating To Billing Aliases", 31 Pages
 Office Action for U.S. Appl. No. 15/465,981, dated Oct. 6, 2021, Kramme, "Identifying Fraudulent Instruments and Identification", 35 Pages.
 Office Action for U.S. Appl. No. 17/078,744, dated Dec. 27, 2021, Kramme, "Reducing False Positives Using Customer Data and Machine Learning", 28 Pages.
 Office Action for U.S. Appl. No. 15/465,838, dated Feb. 1, 2022, Kramme, "Detecting Financial Fraud Using Spending Pattern Data for Individual Cardholders", 41 Pages.
 Office Action for U.S. Appl. No. 17/134,901, dated Feb. 18, 2022, Kramme, "Reducing False Positive Fraud Alerts for Online Financial Transaction", 10 Pages.

* cited by examiner

10

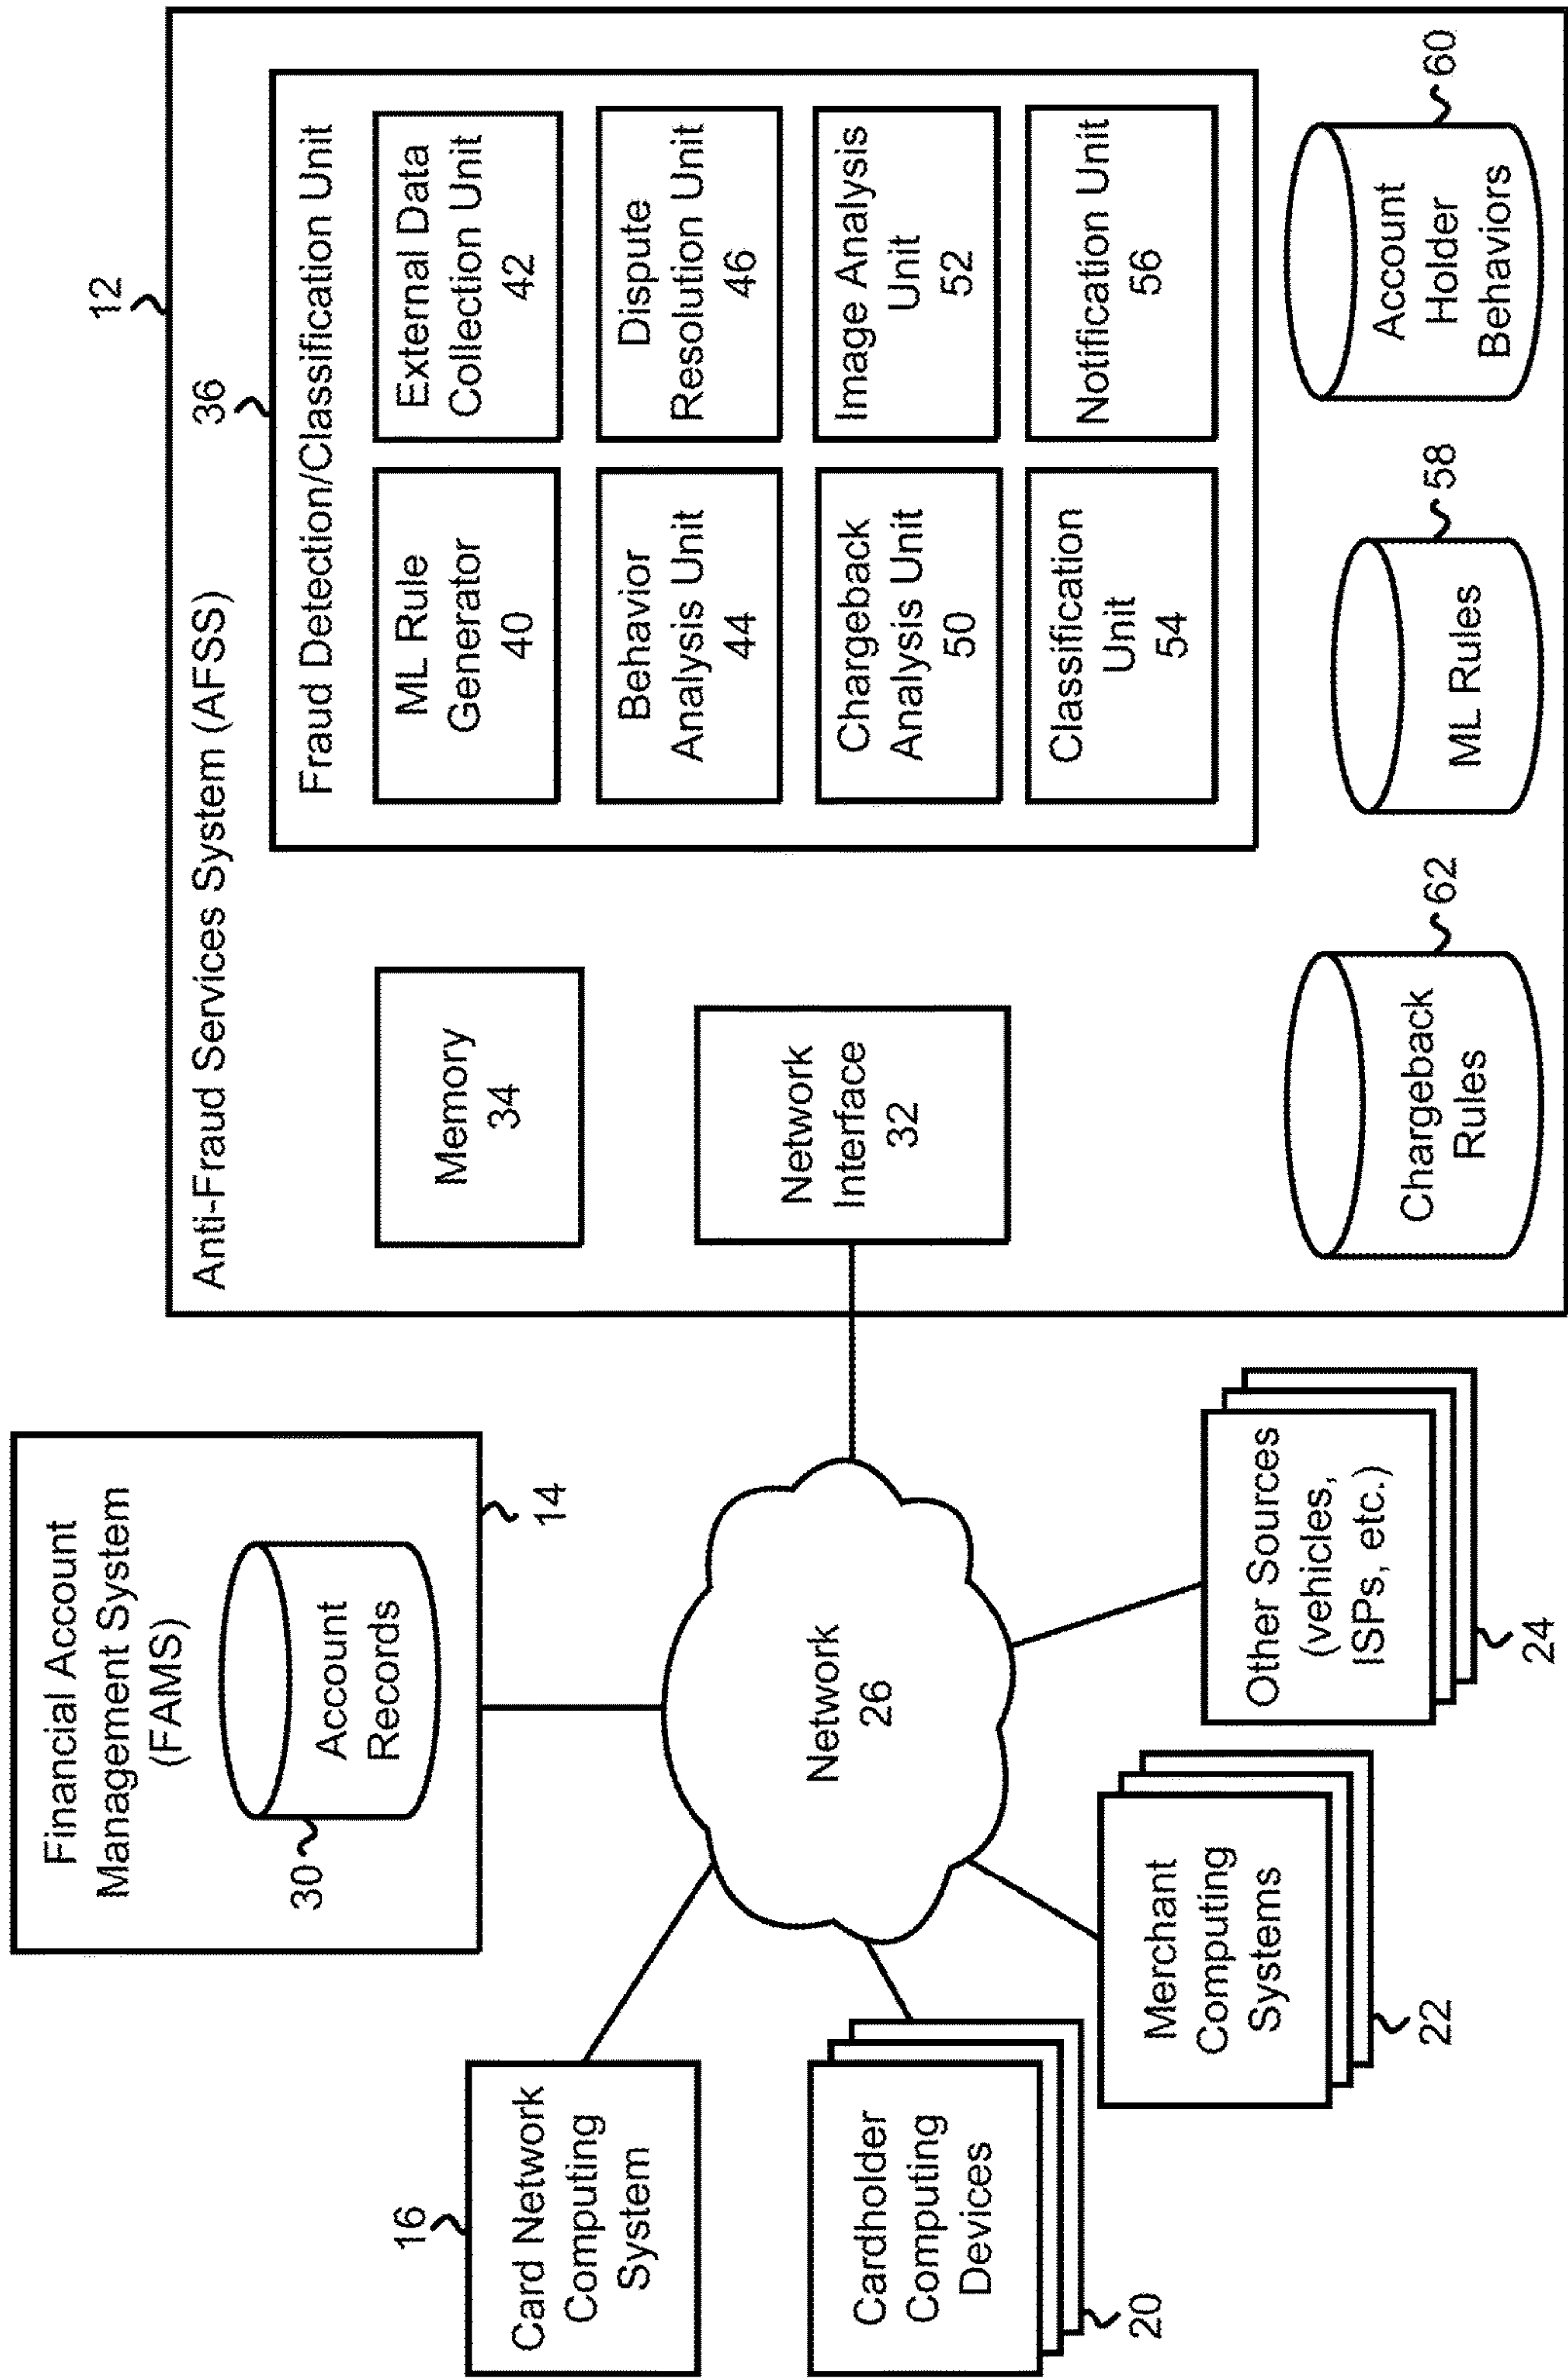


FIG. 1

80

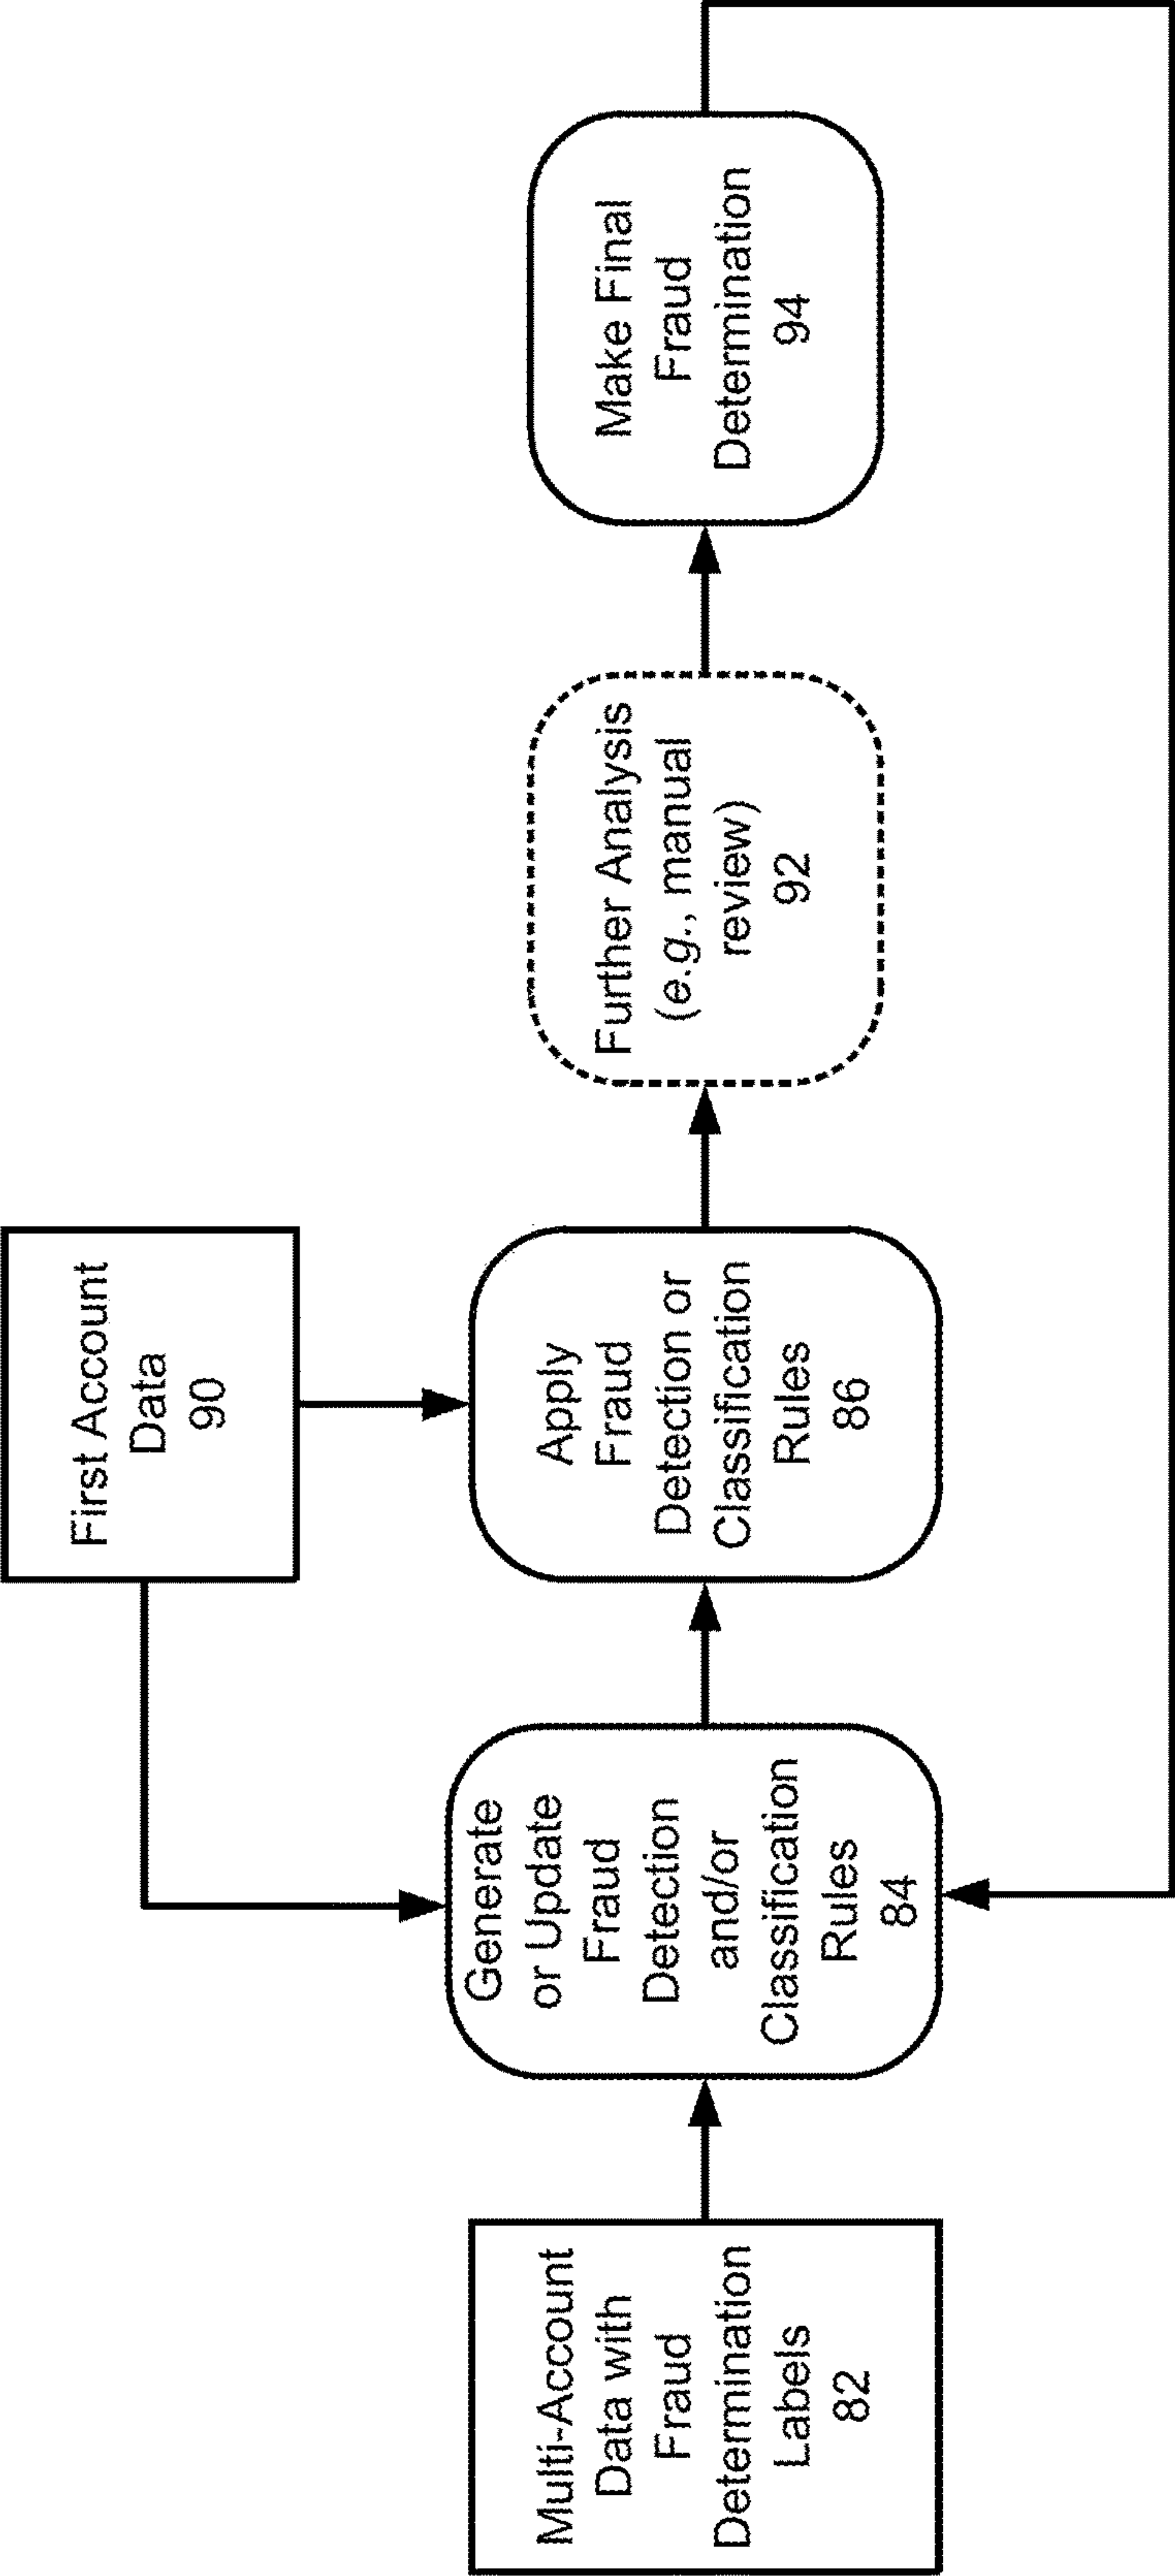


FIG. 2

100

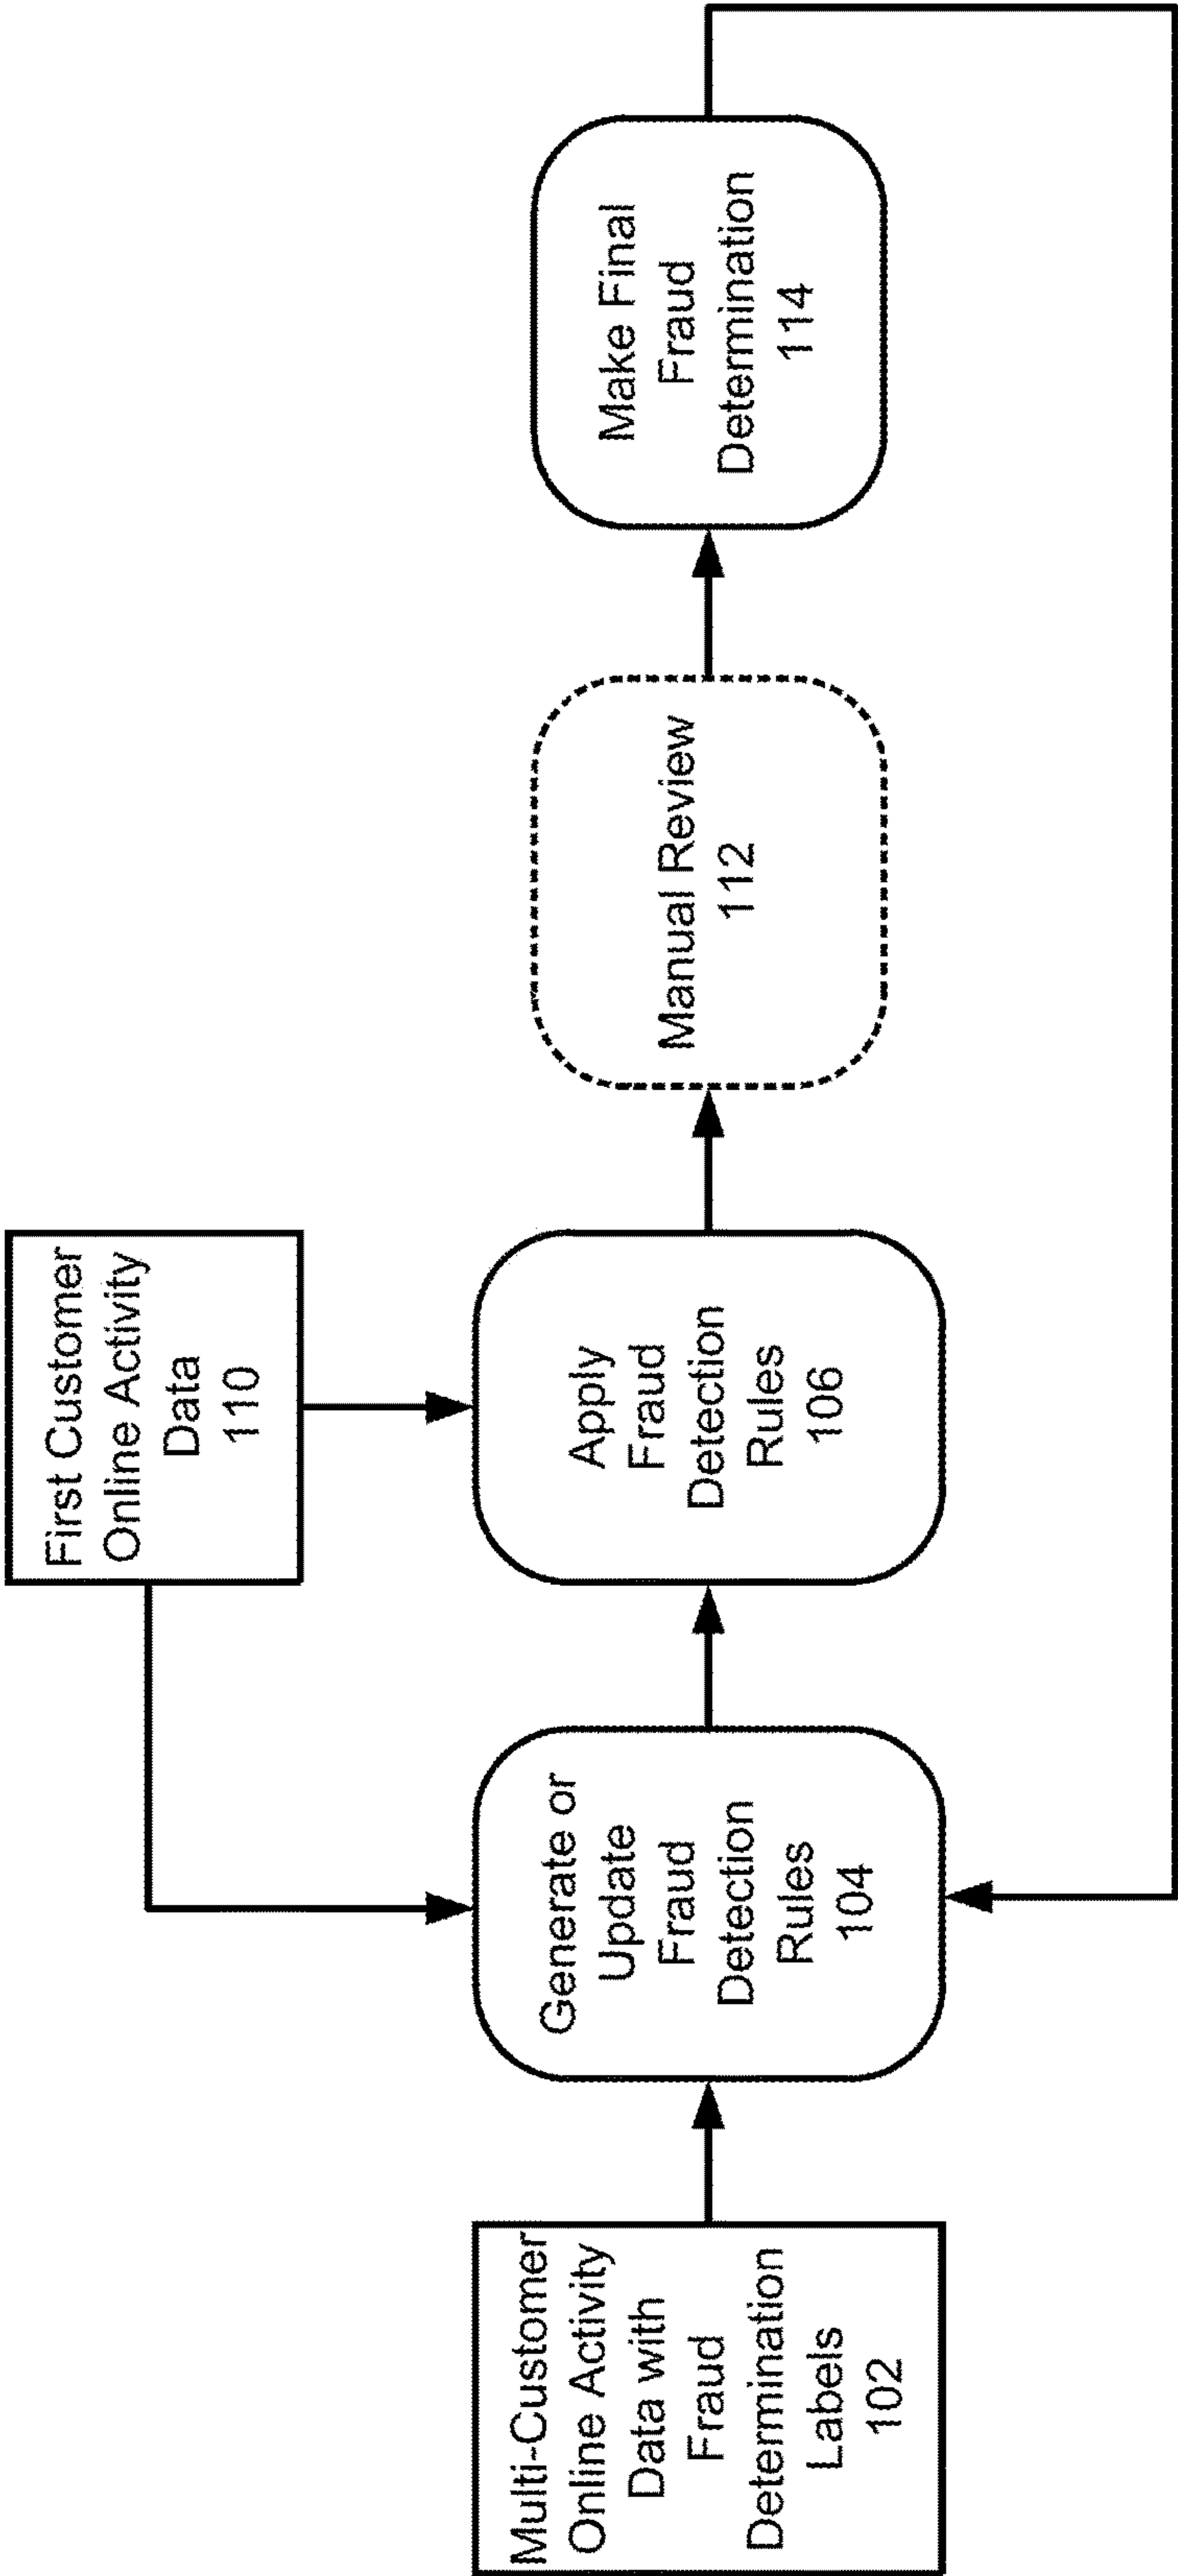


FIG. 3A

120

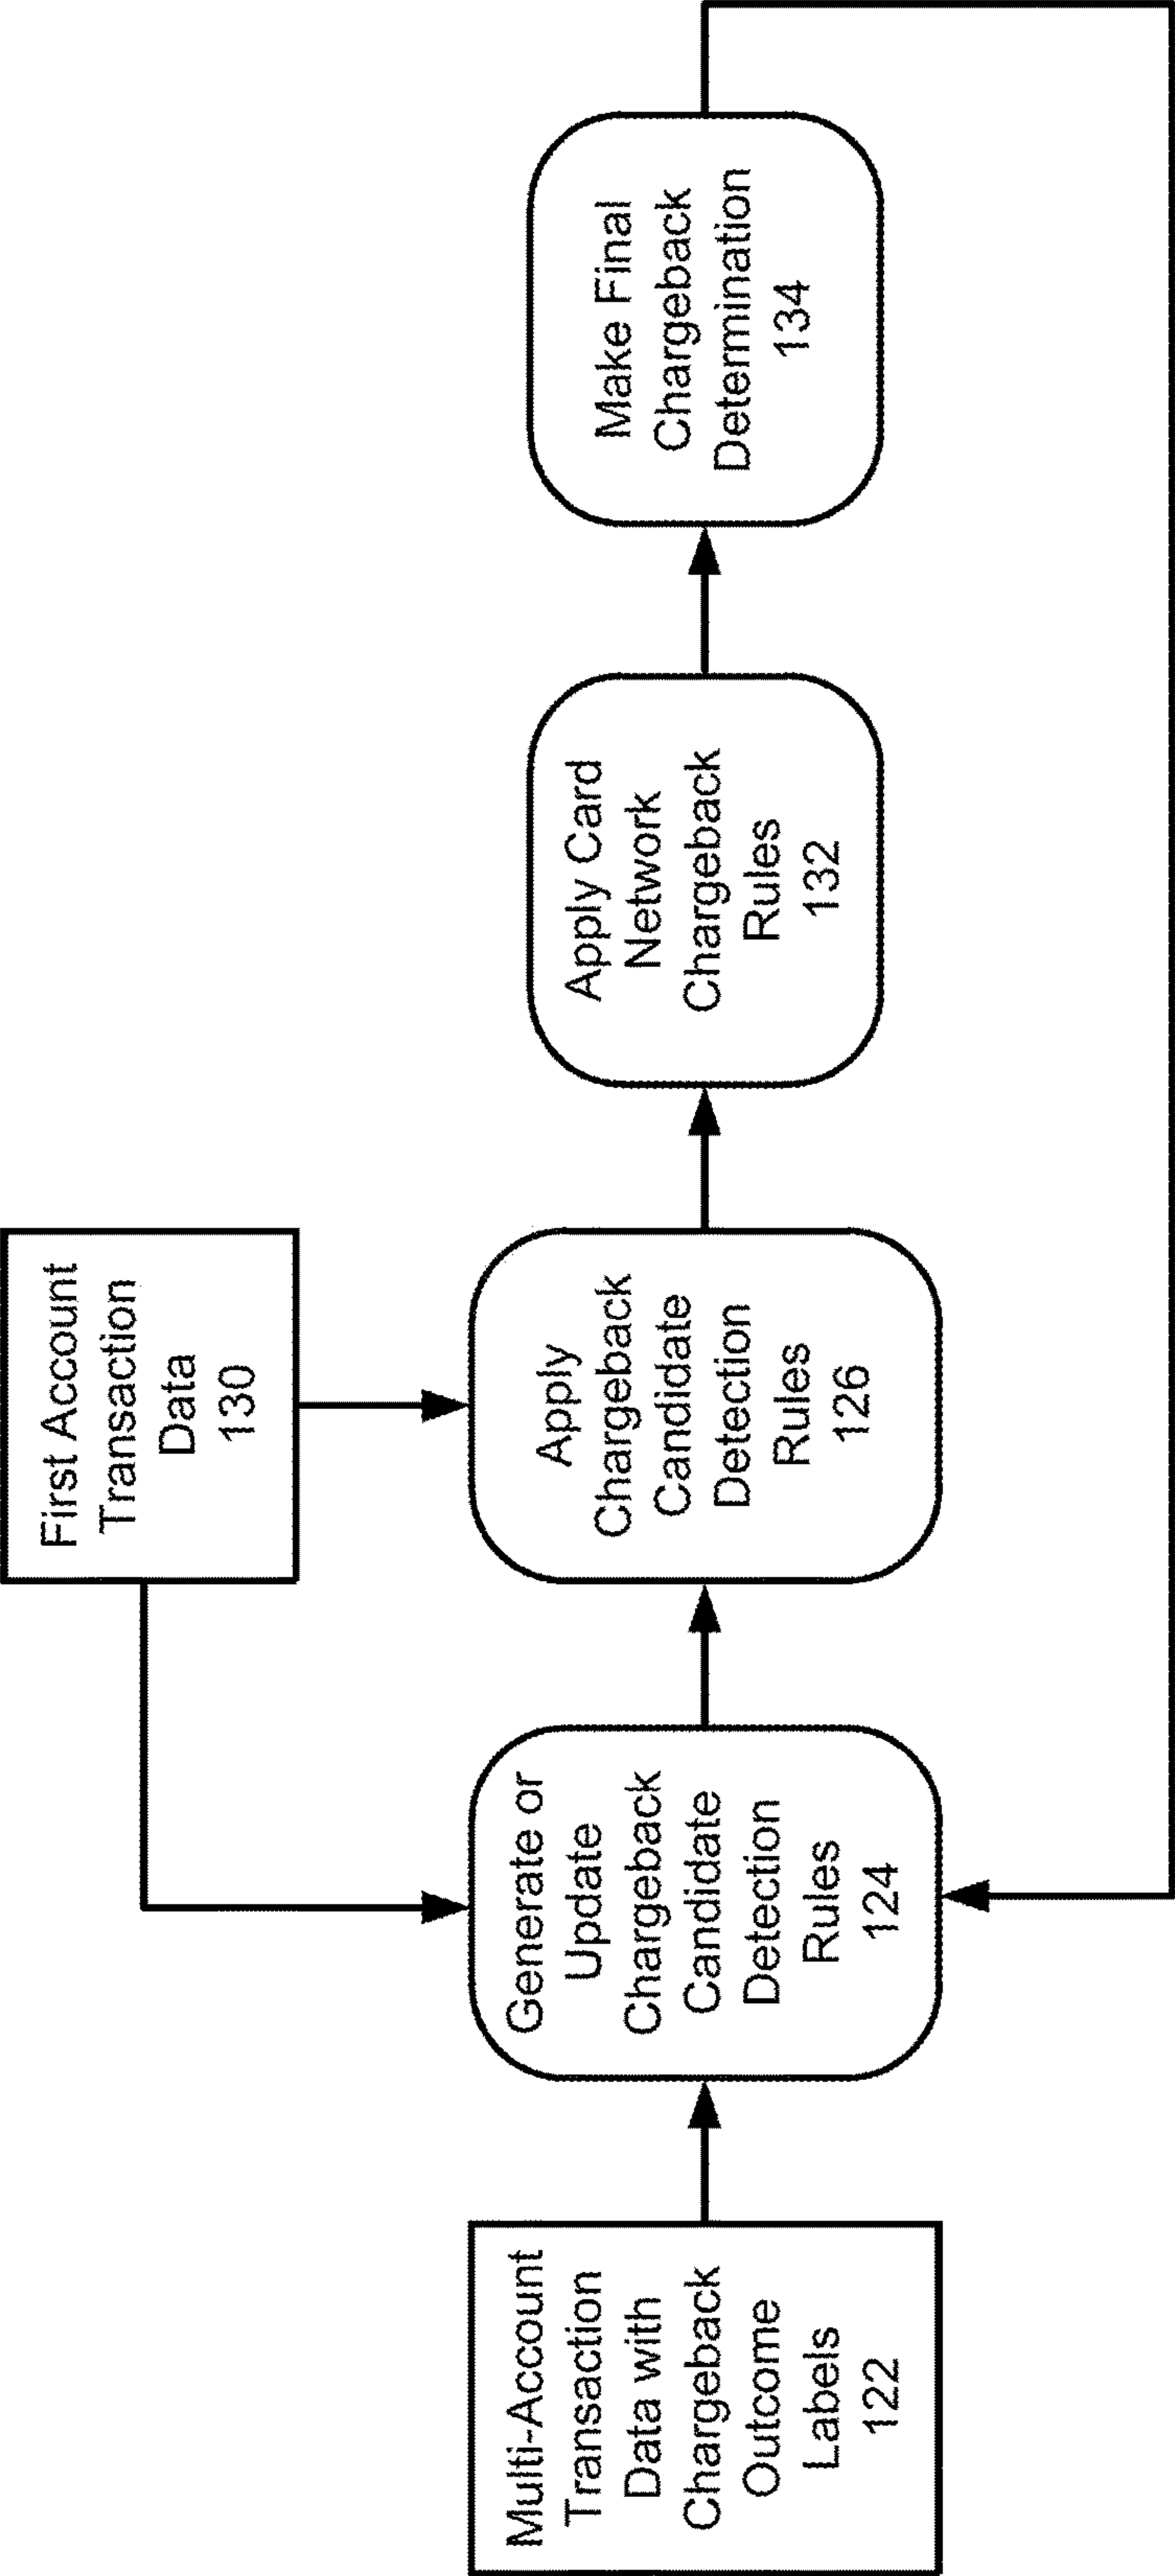


FIG. 3B

140

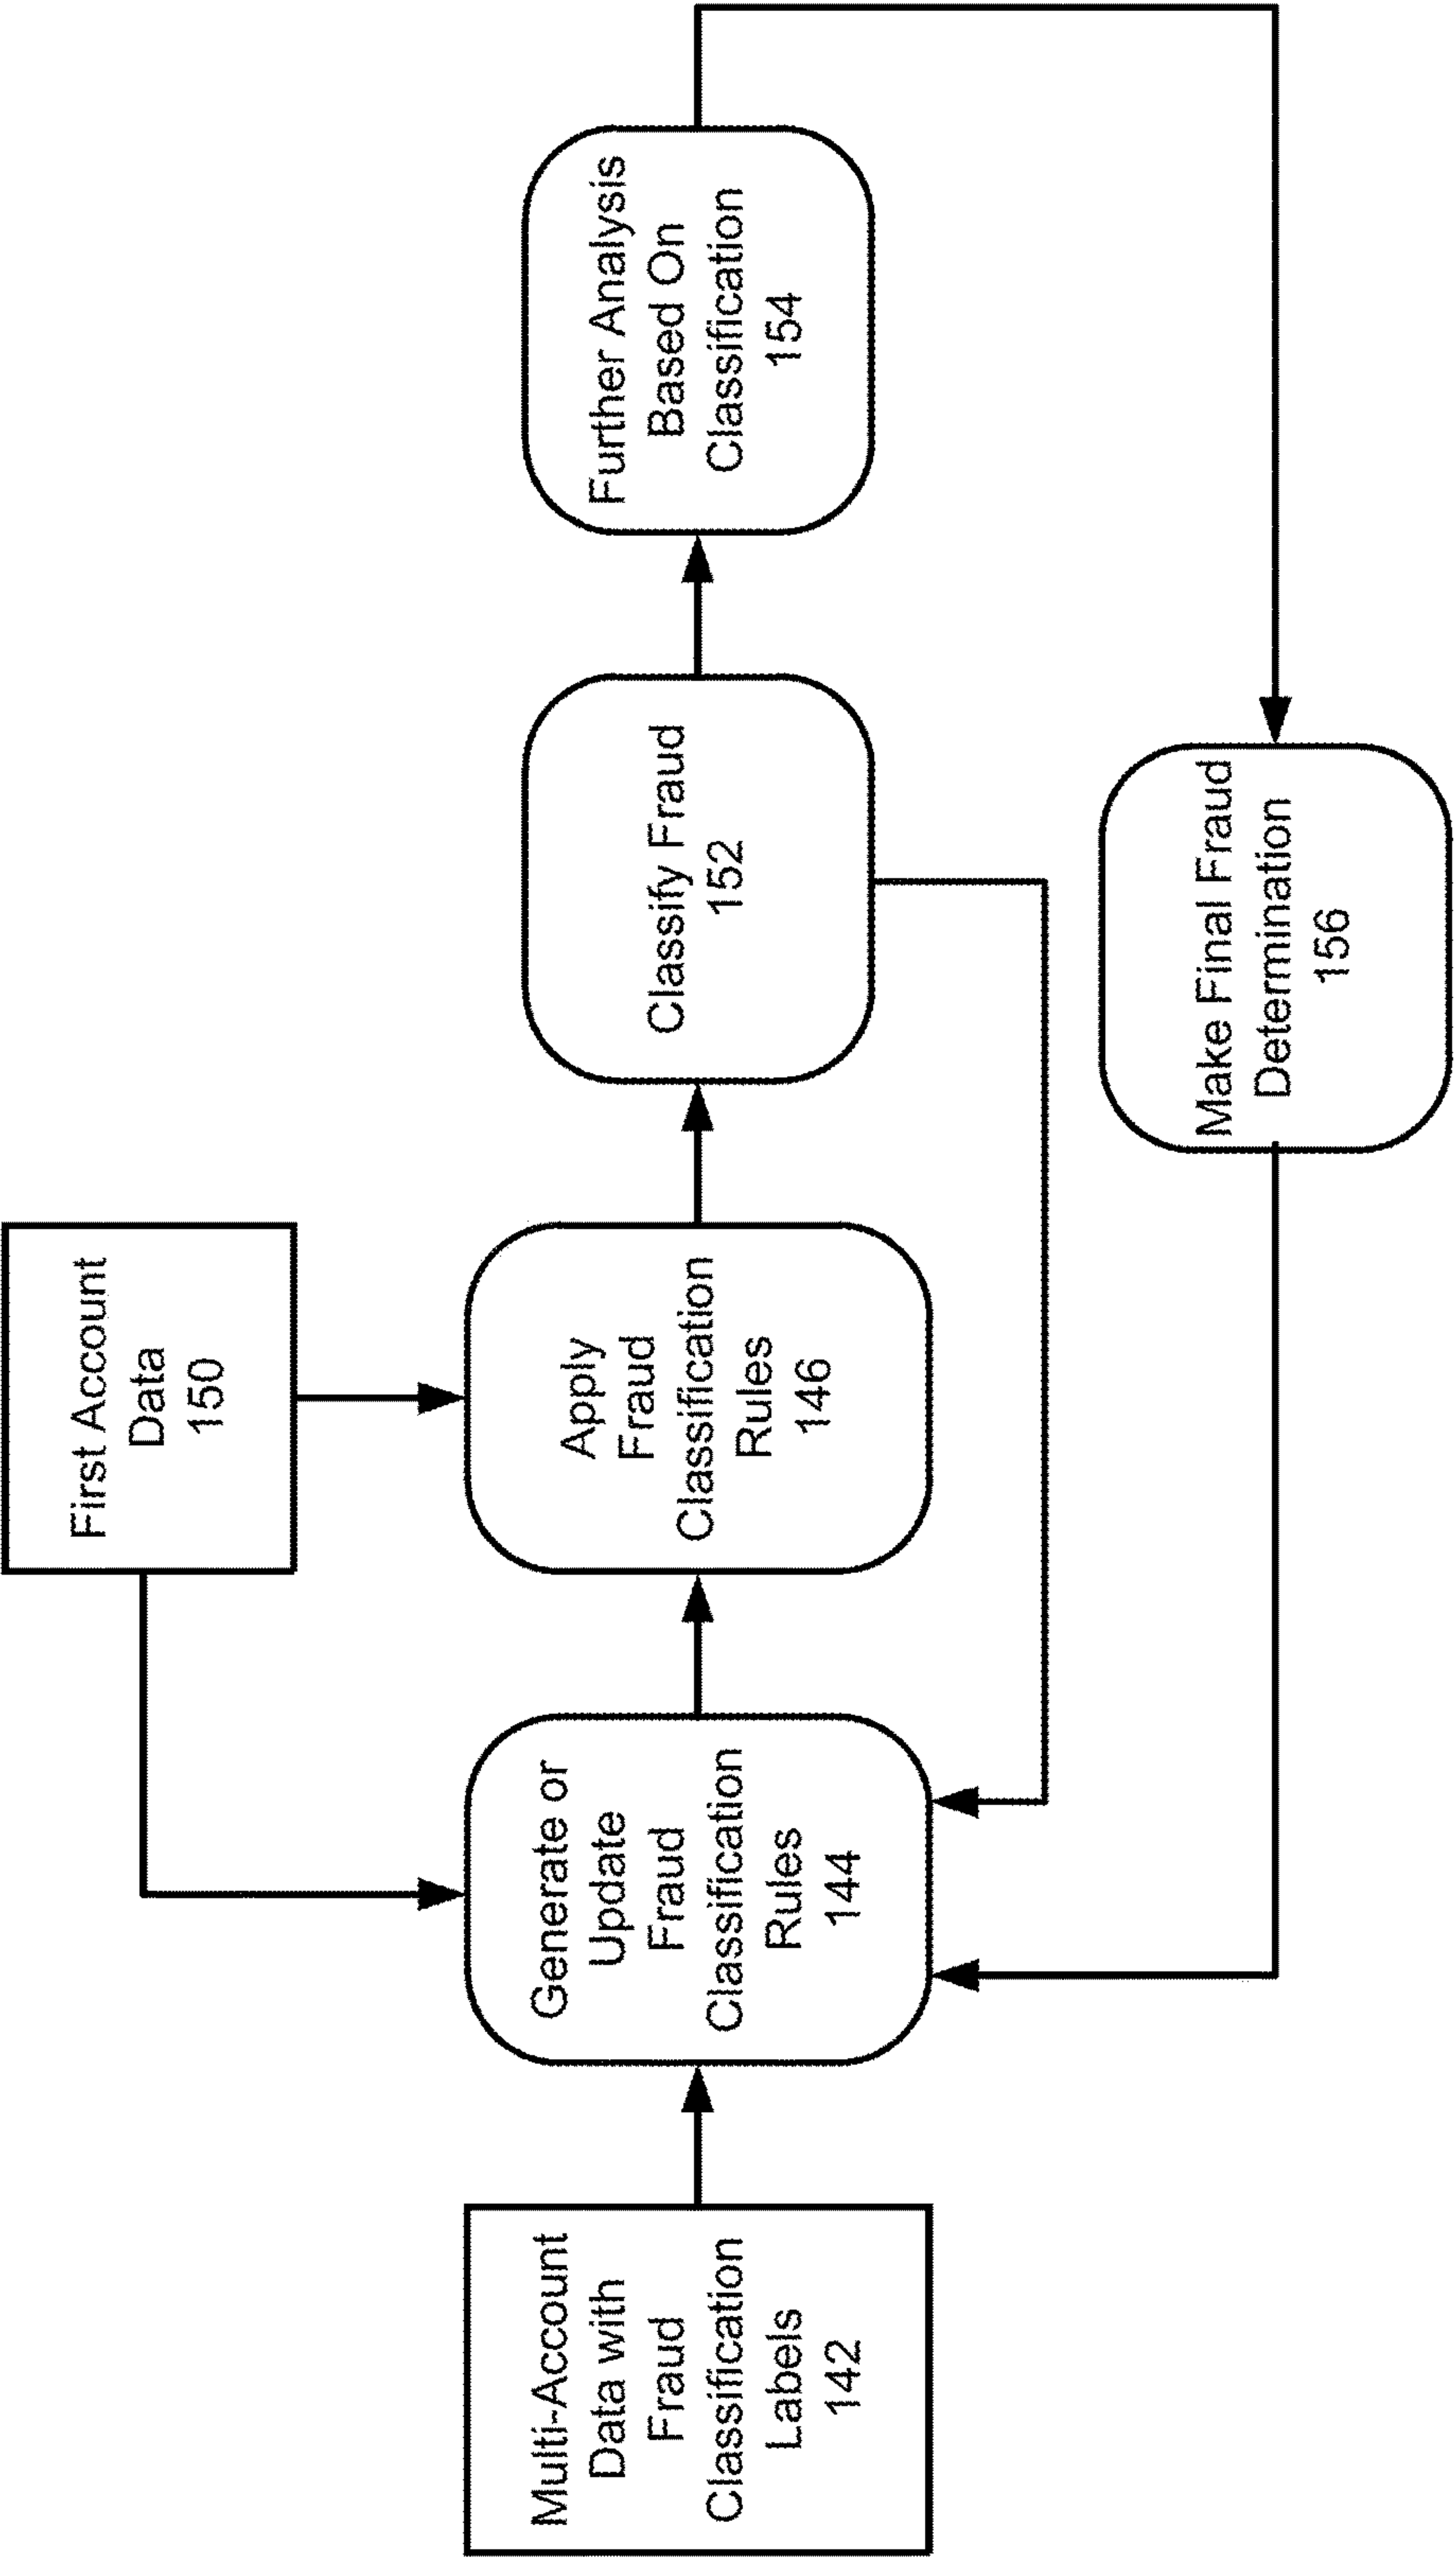


FIG. 3C

160

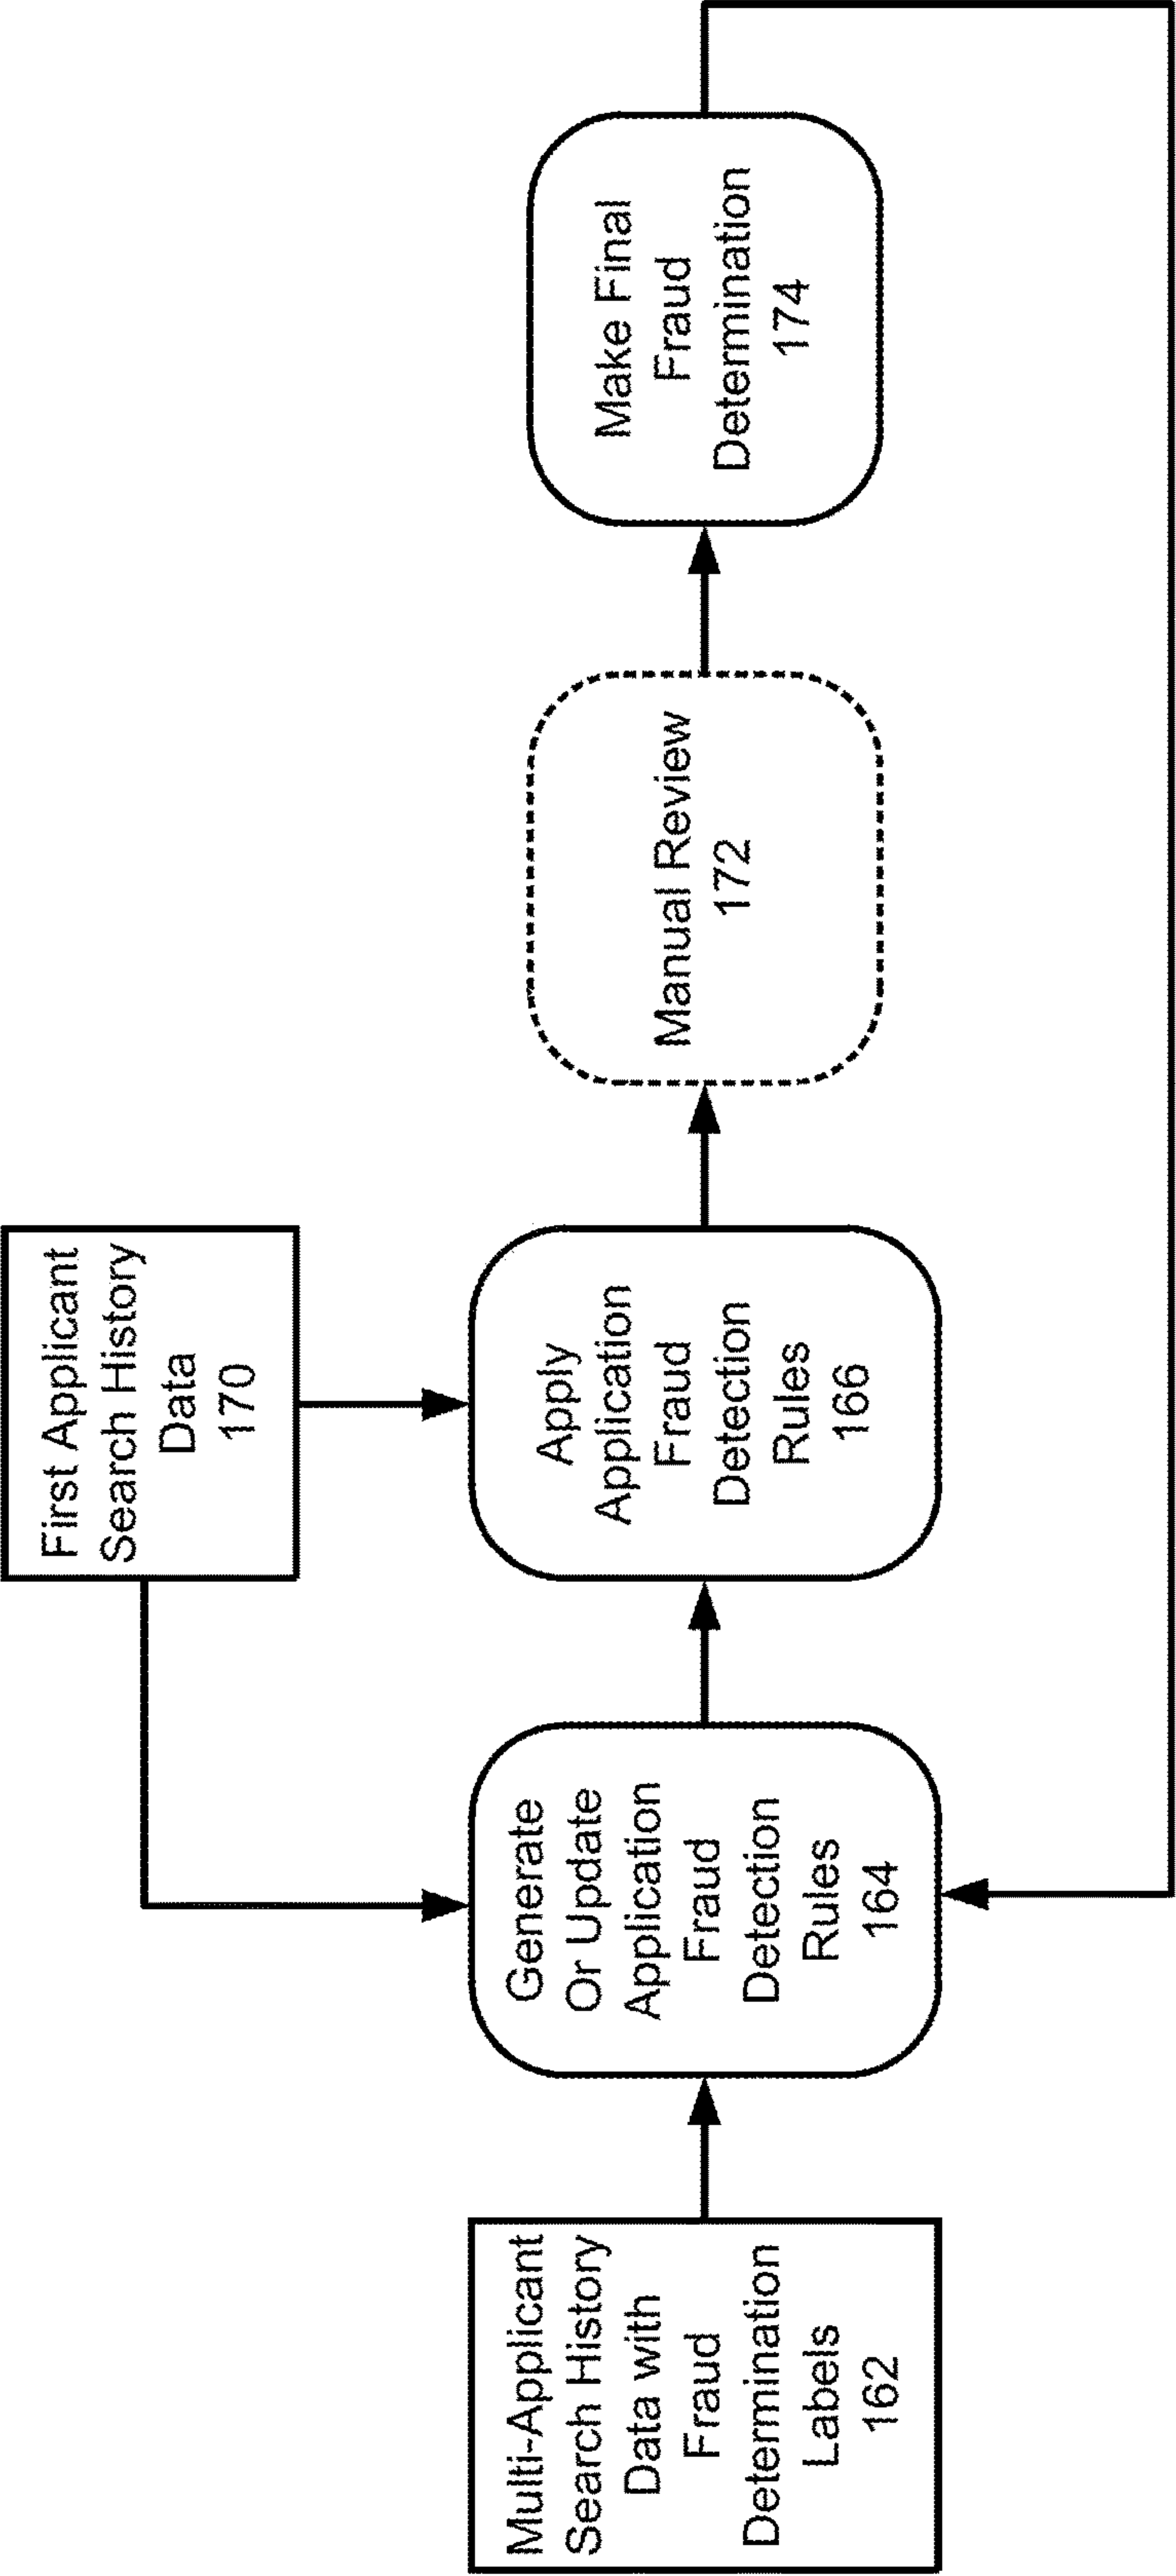


FIG. 3D

180

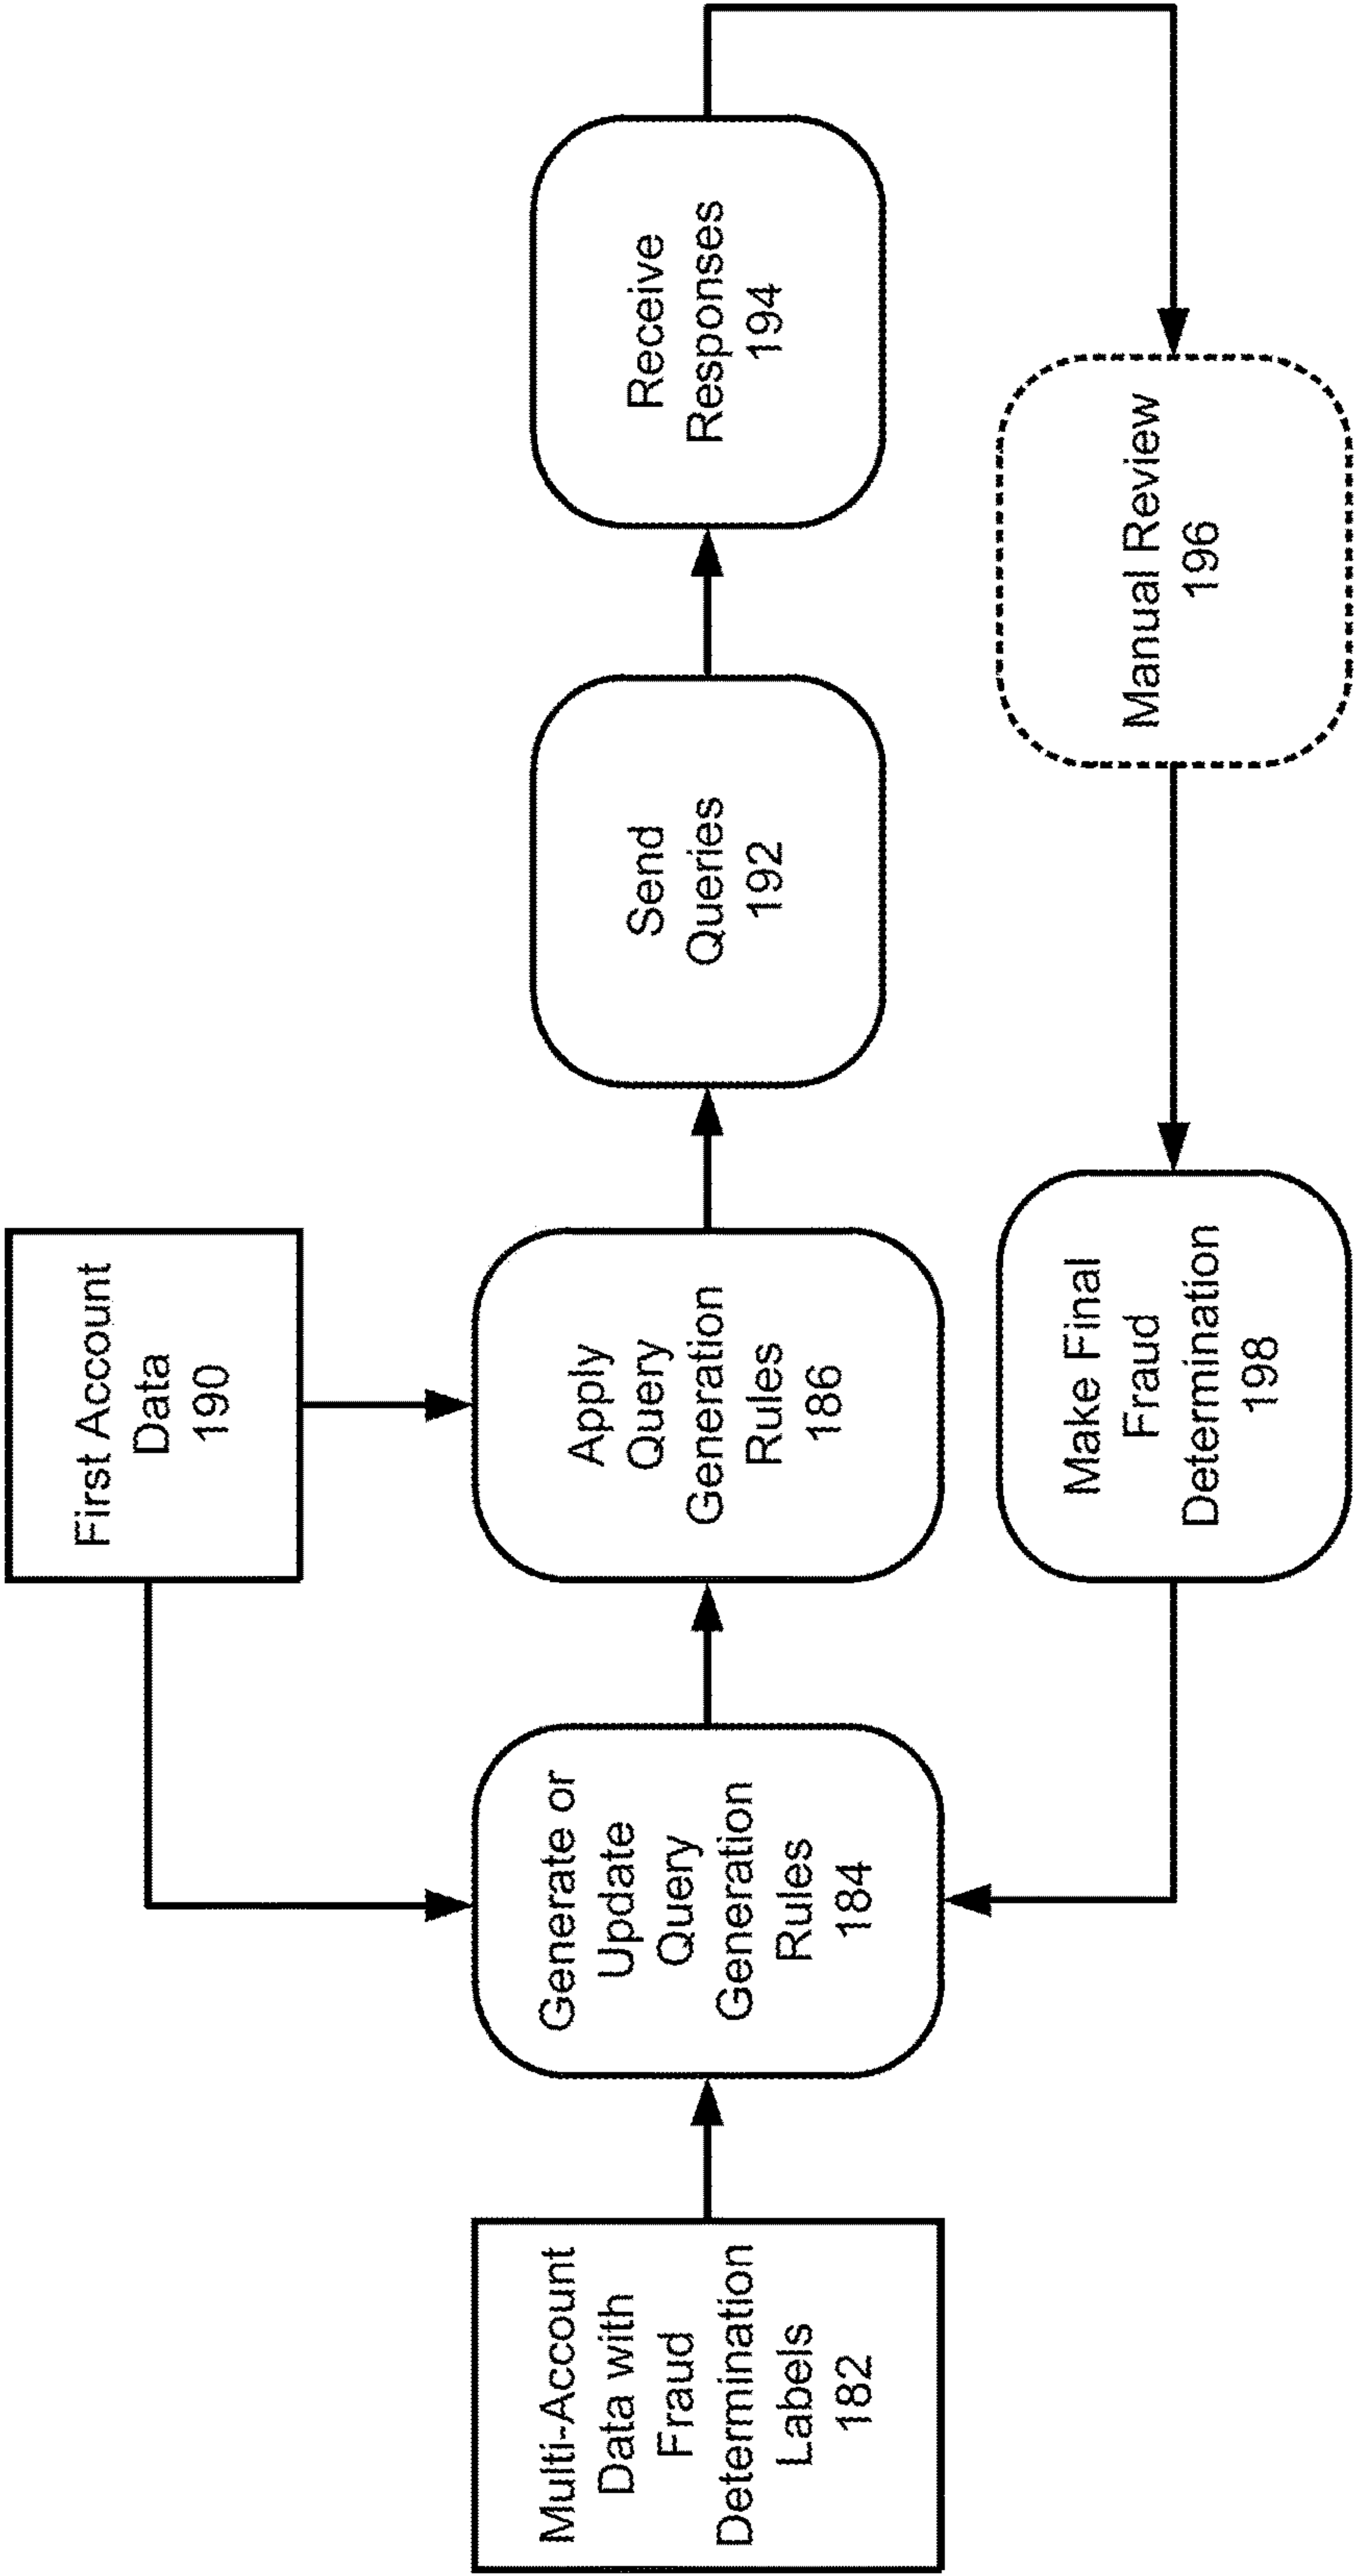


FIG. 3E

200

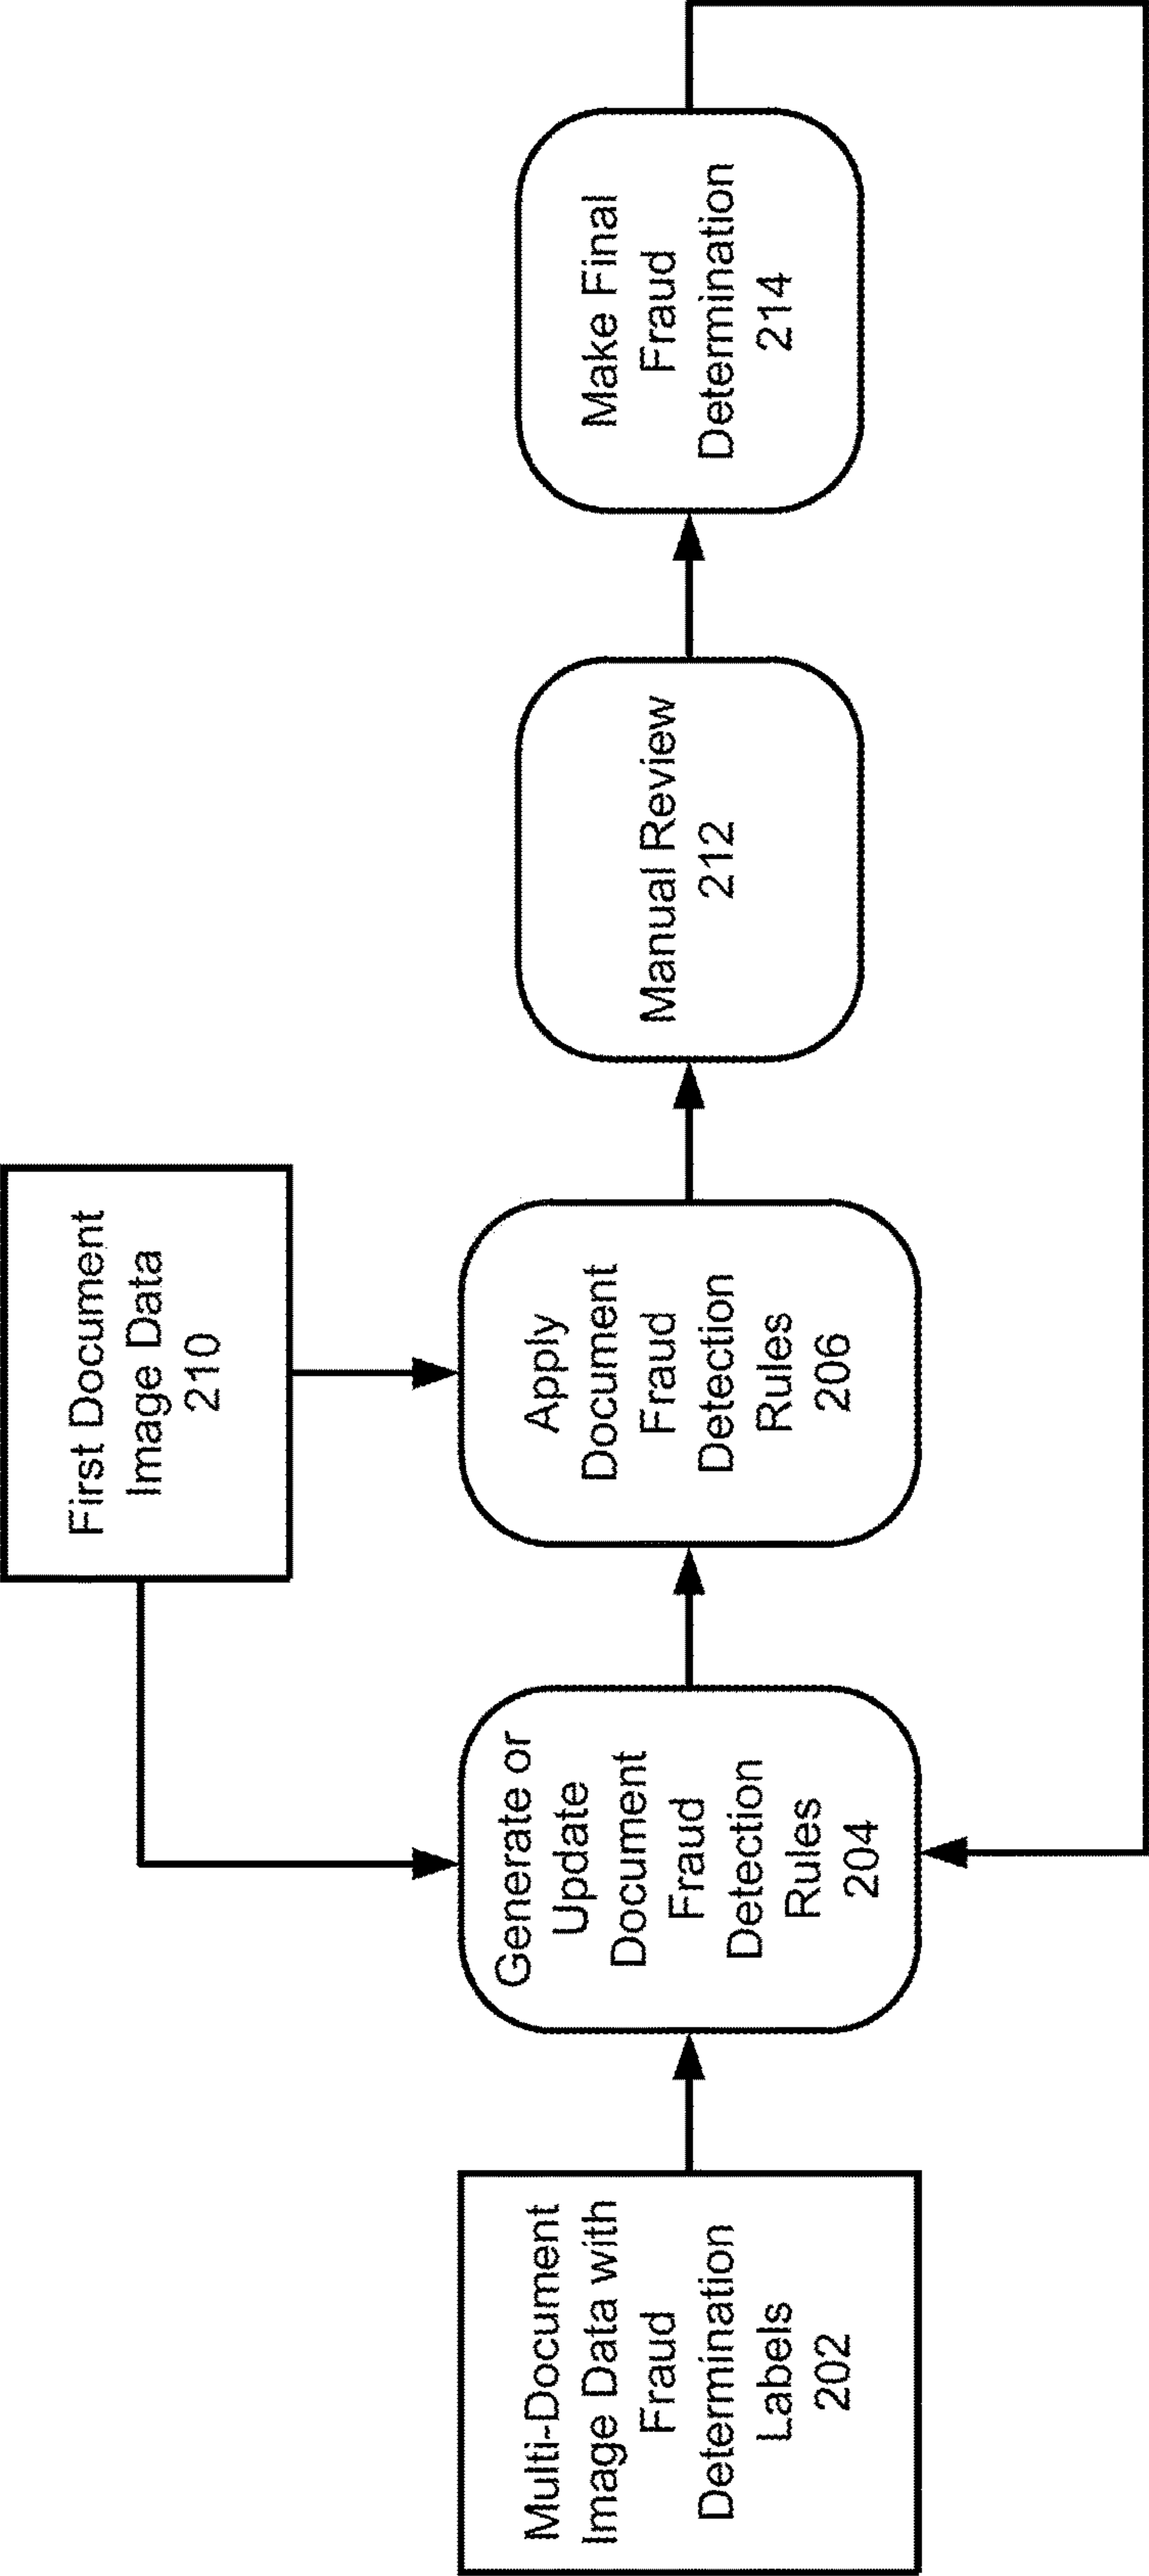
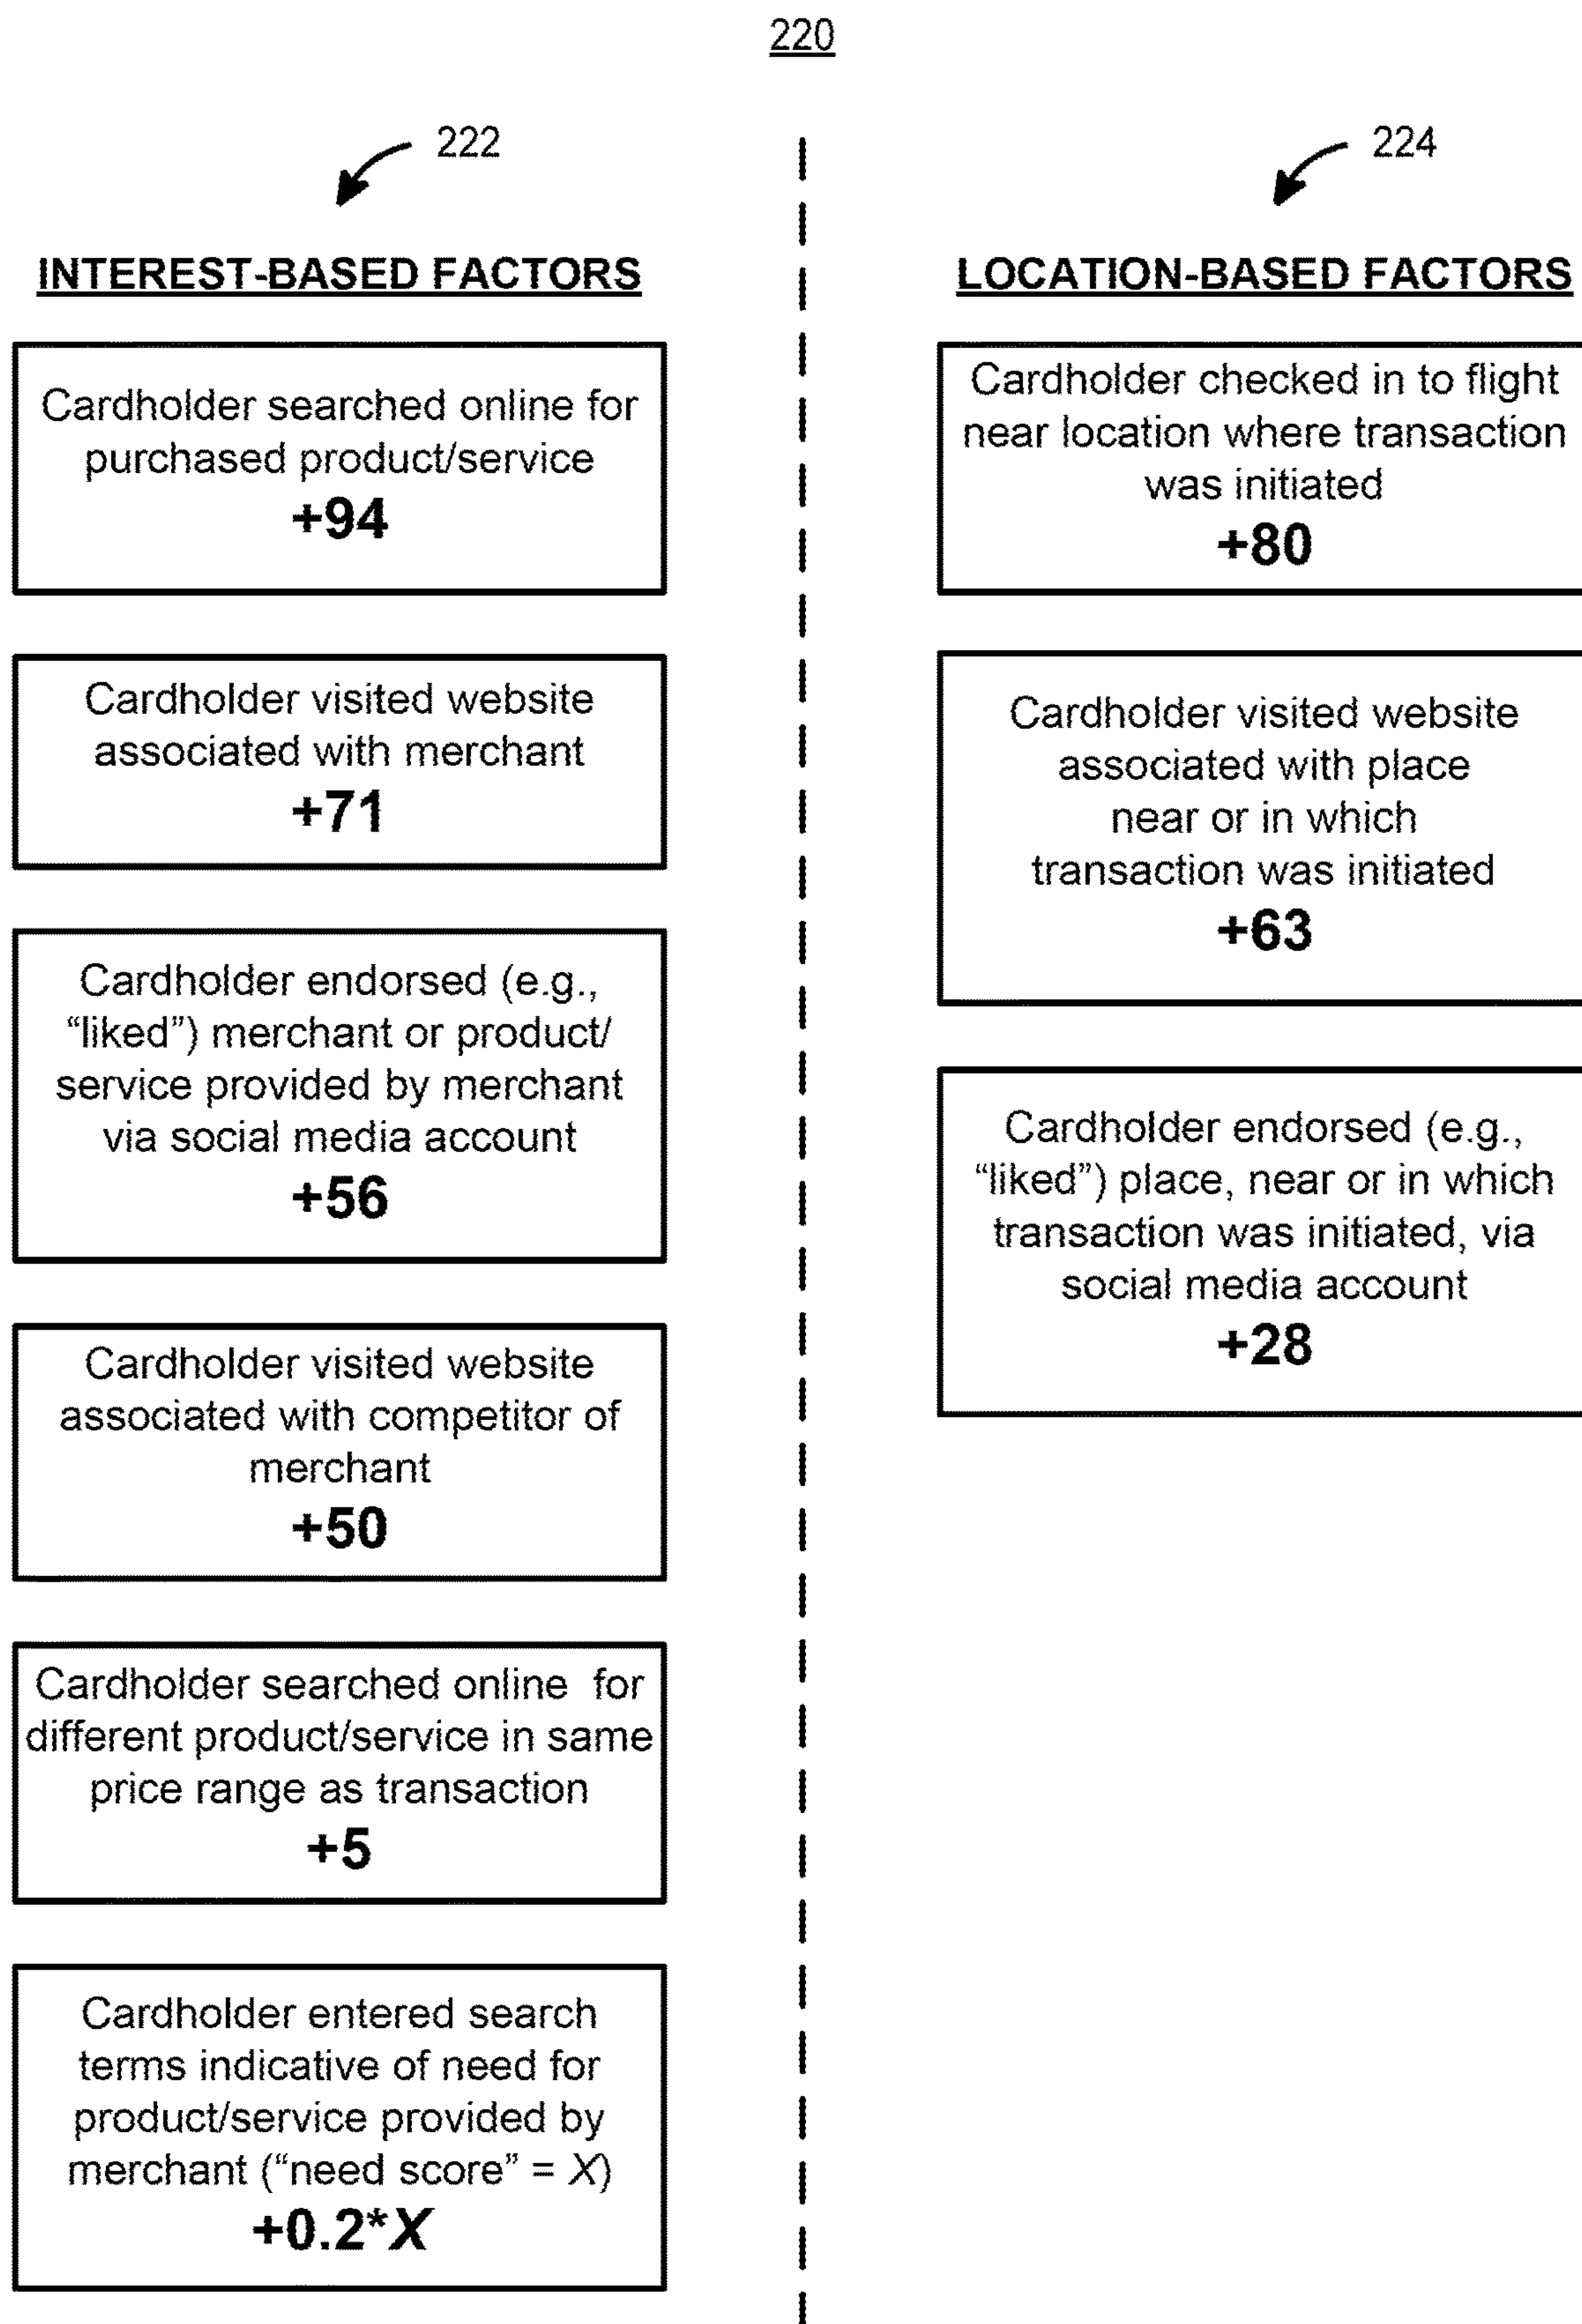


FIG. 3F

**FIG. 4A**

230

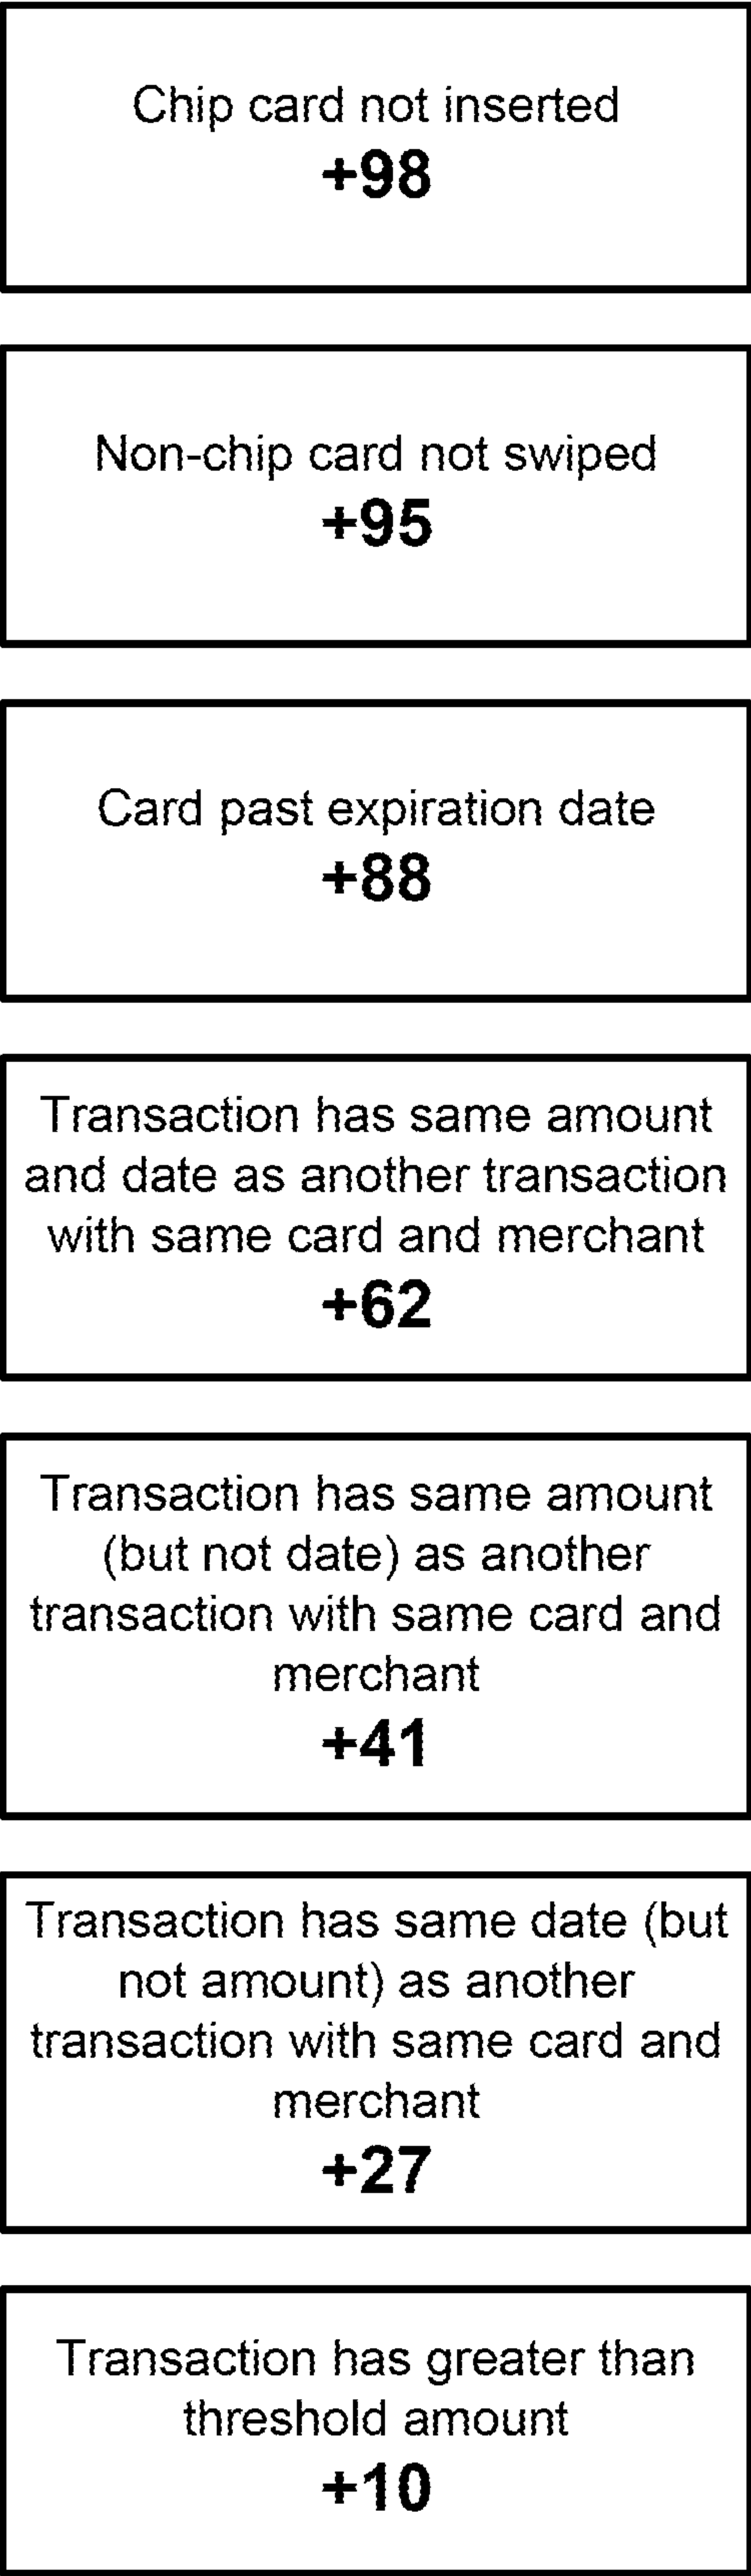


FIG. 4B

240

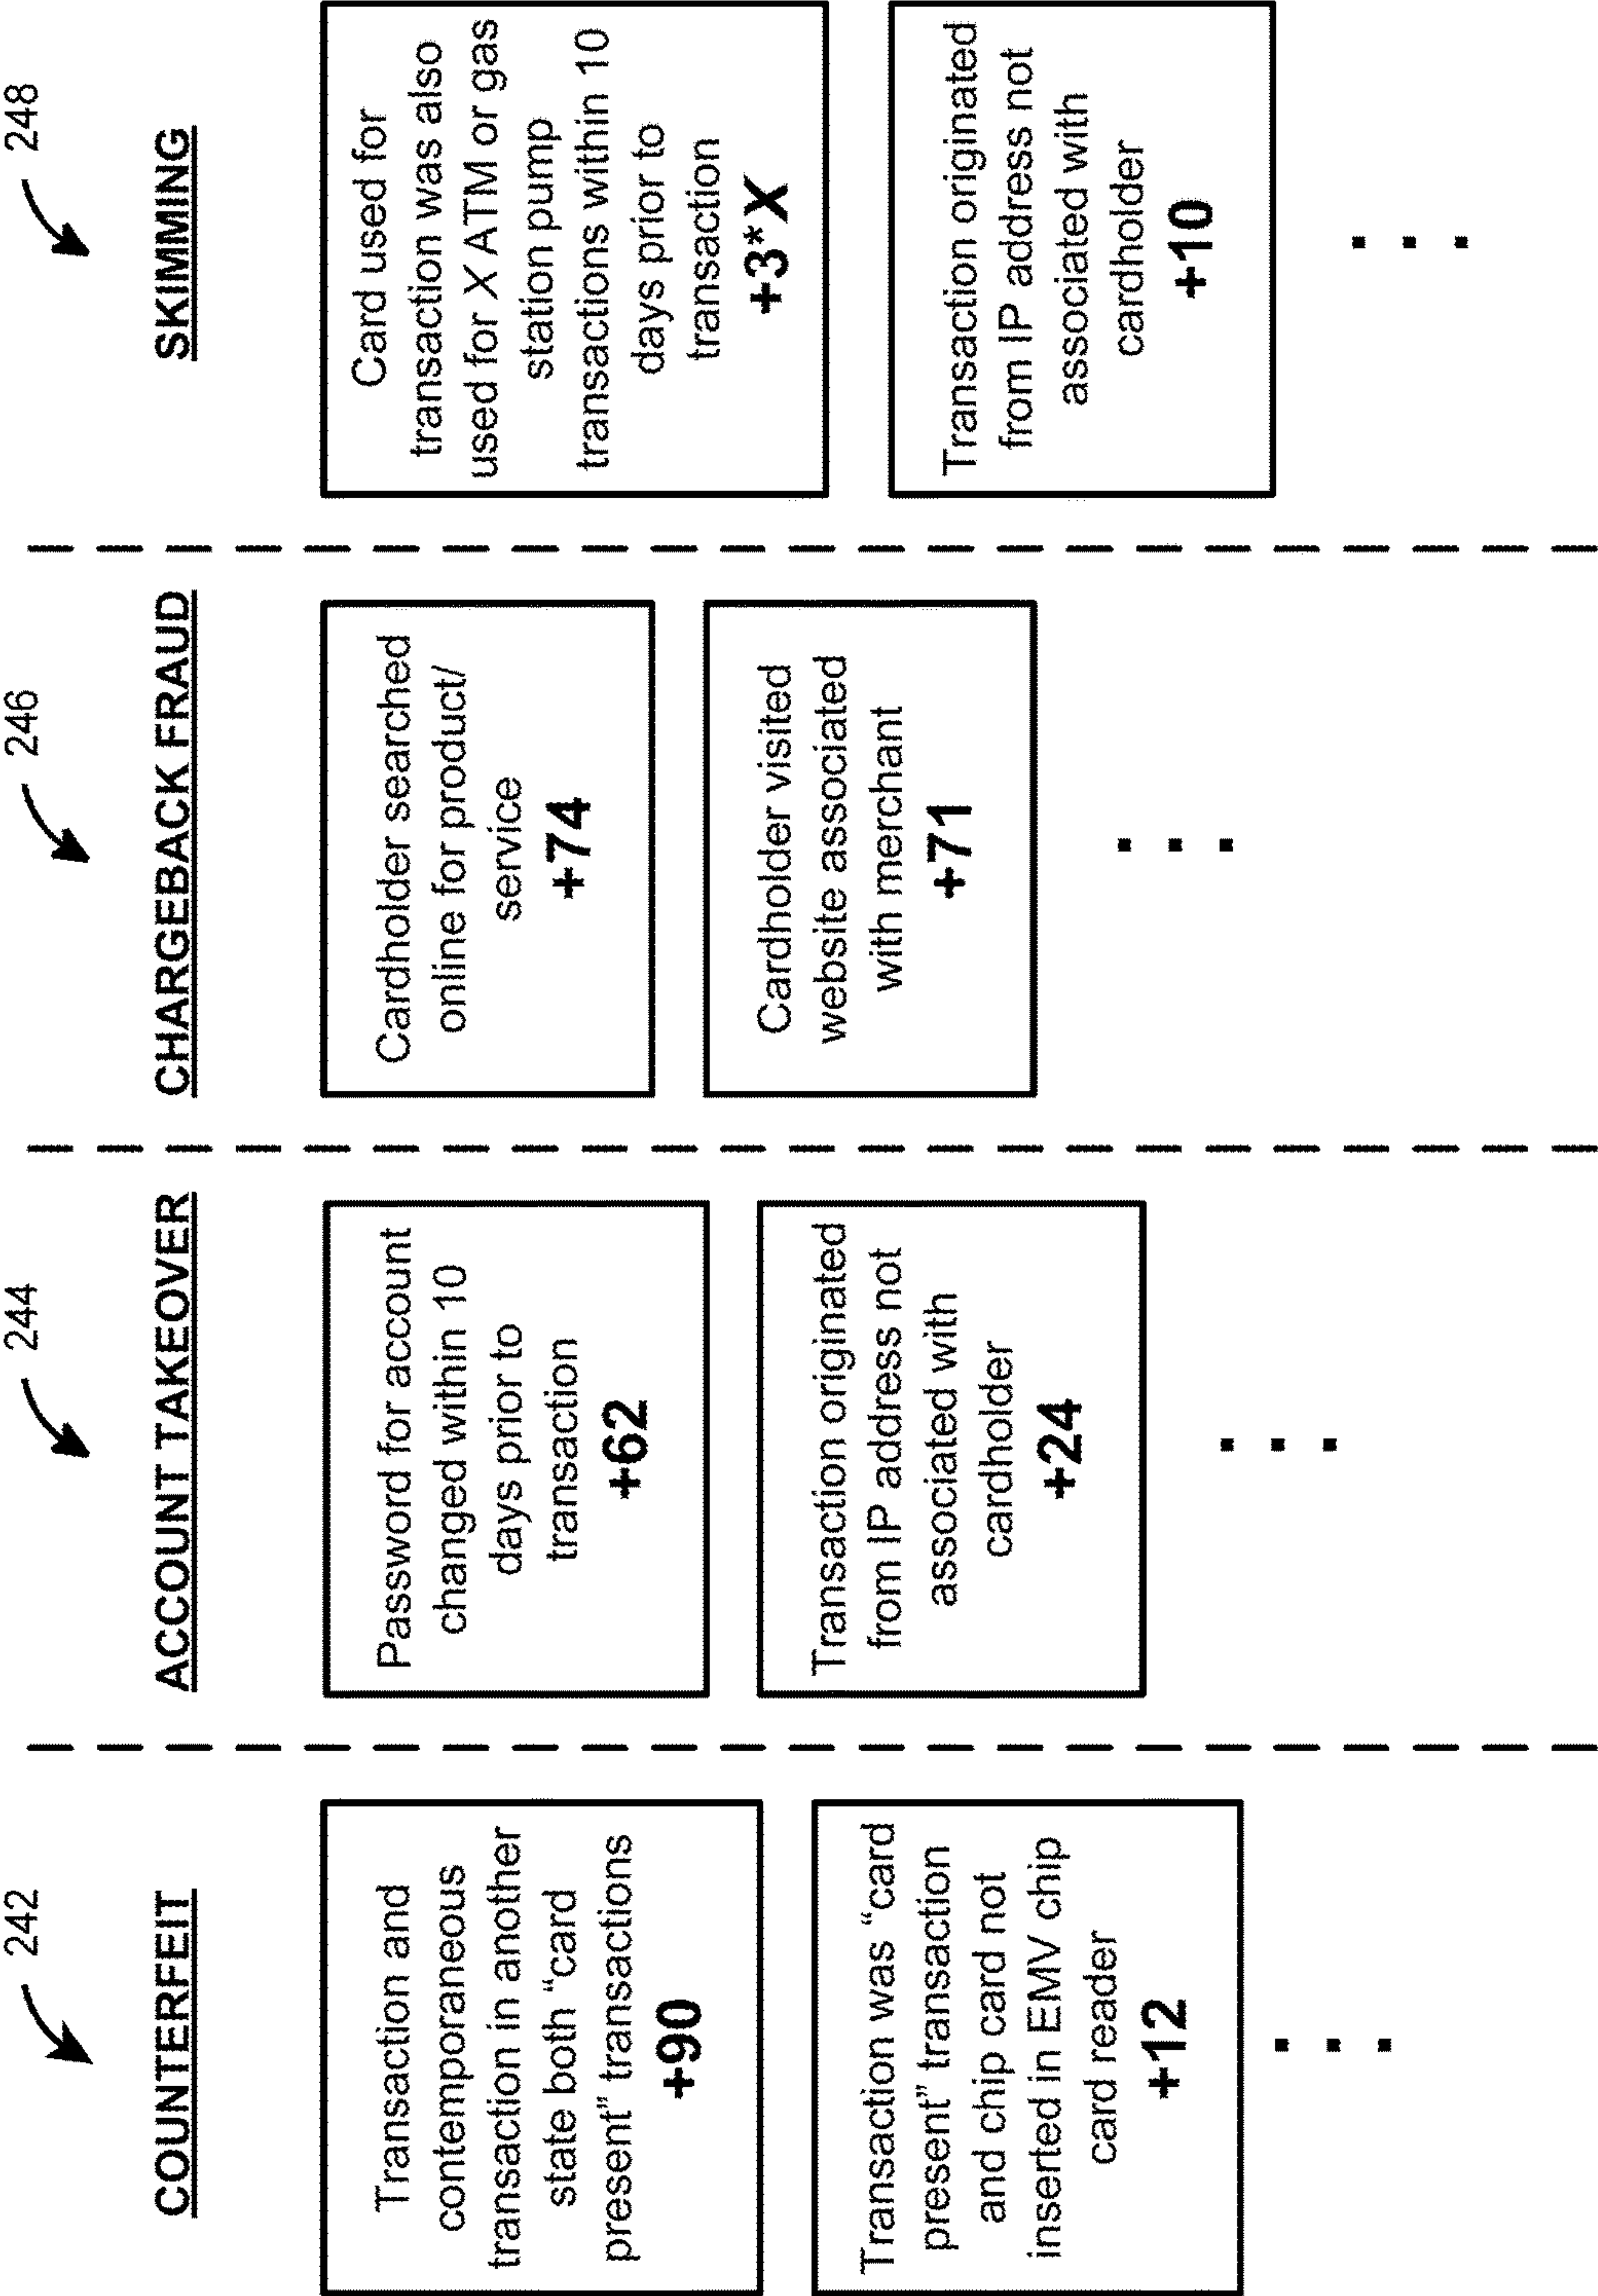


FIG. 4C

260

Search term included name of
purported applicant

+60

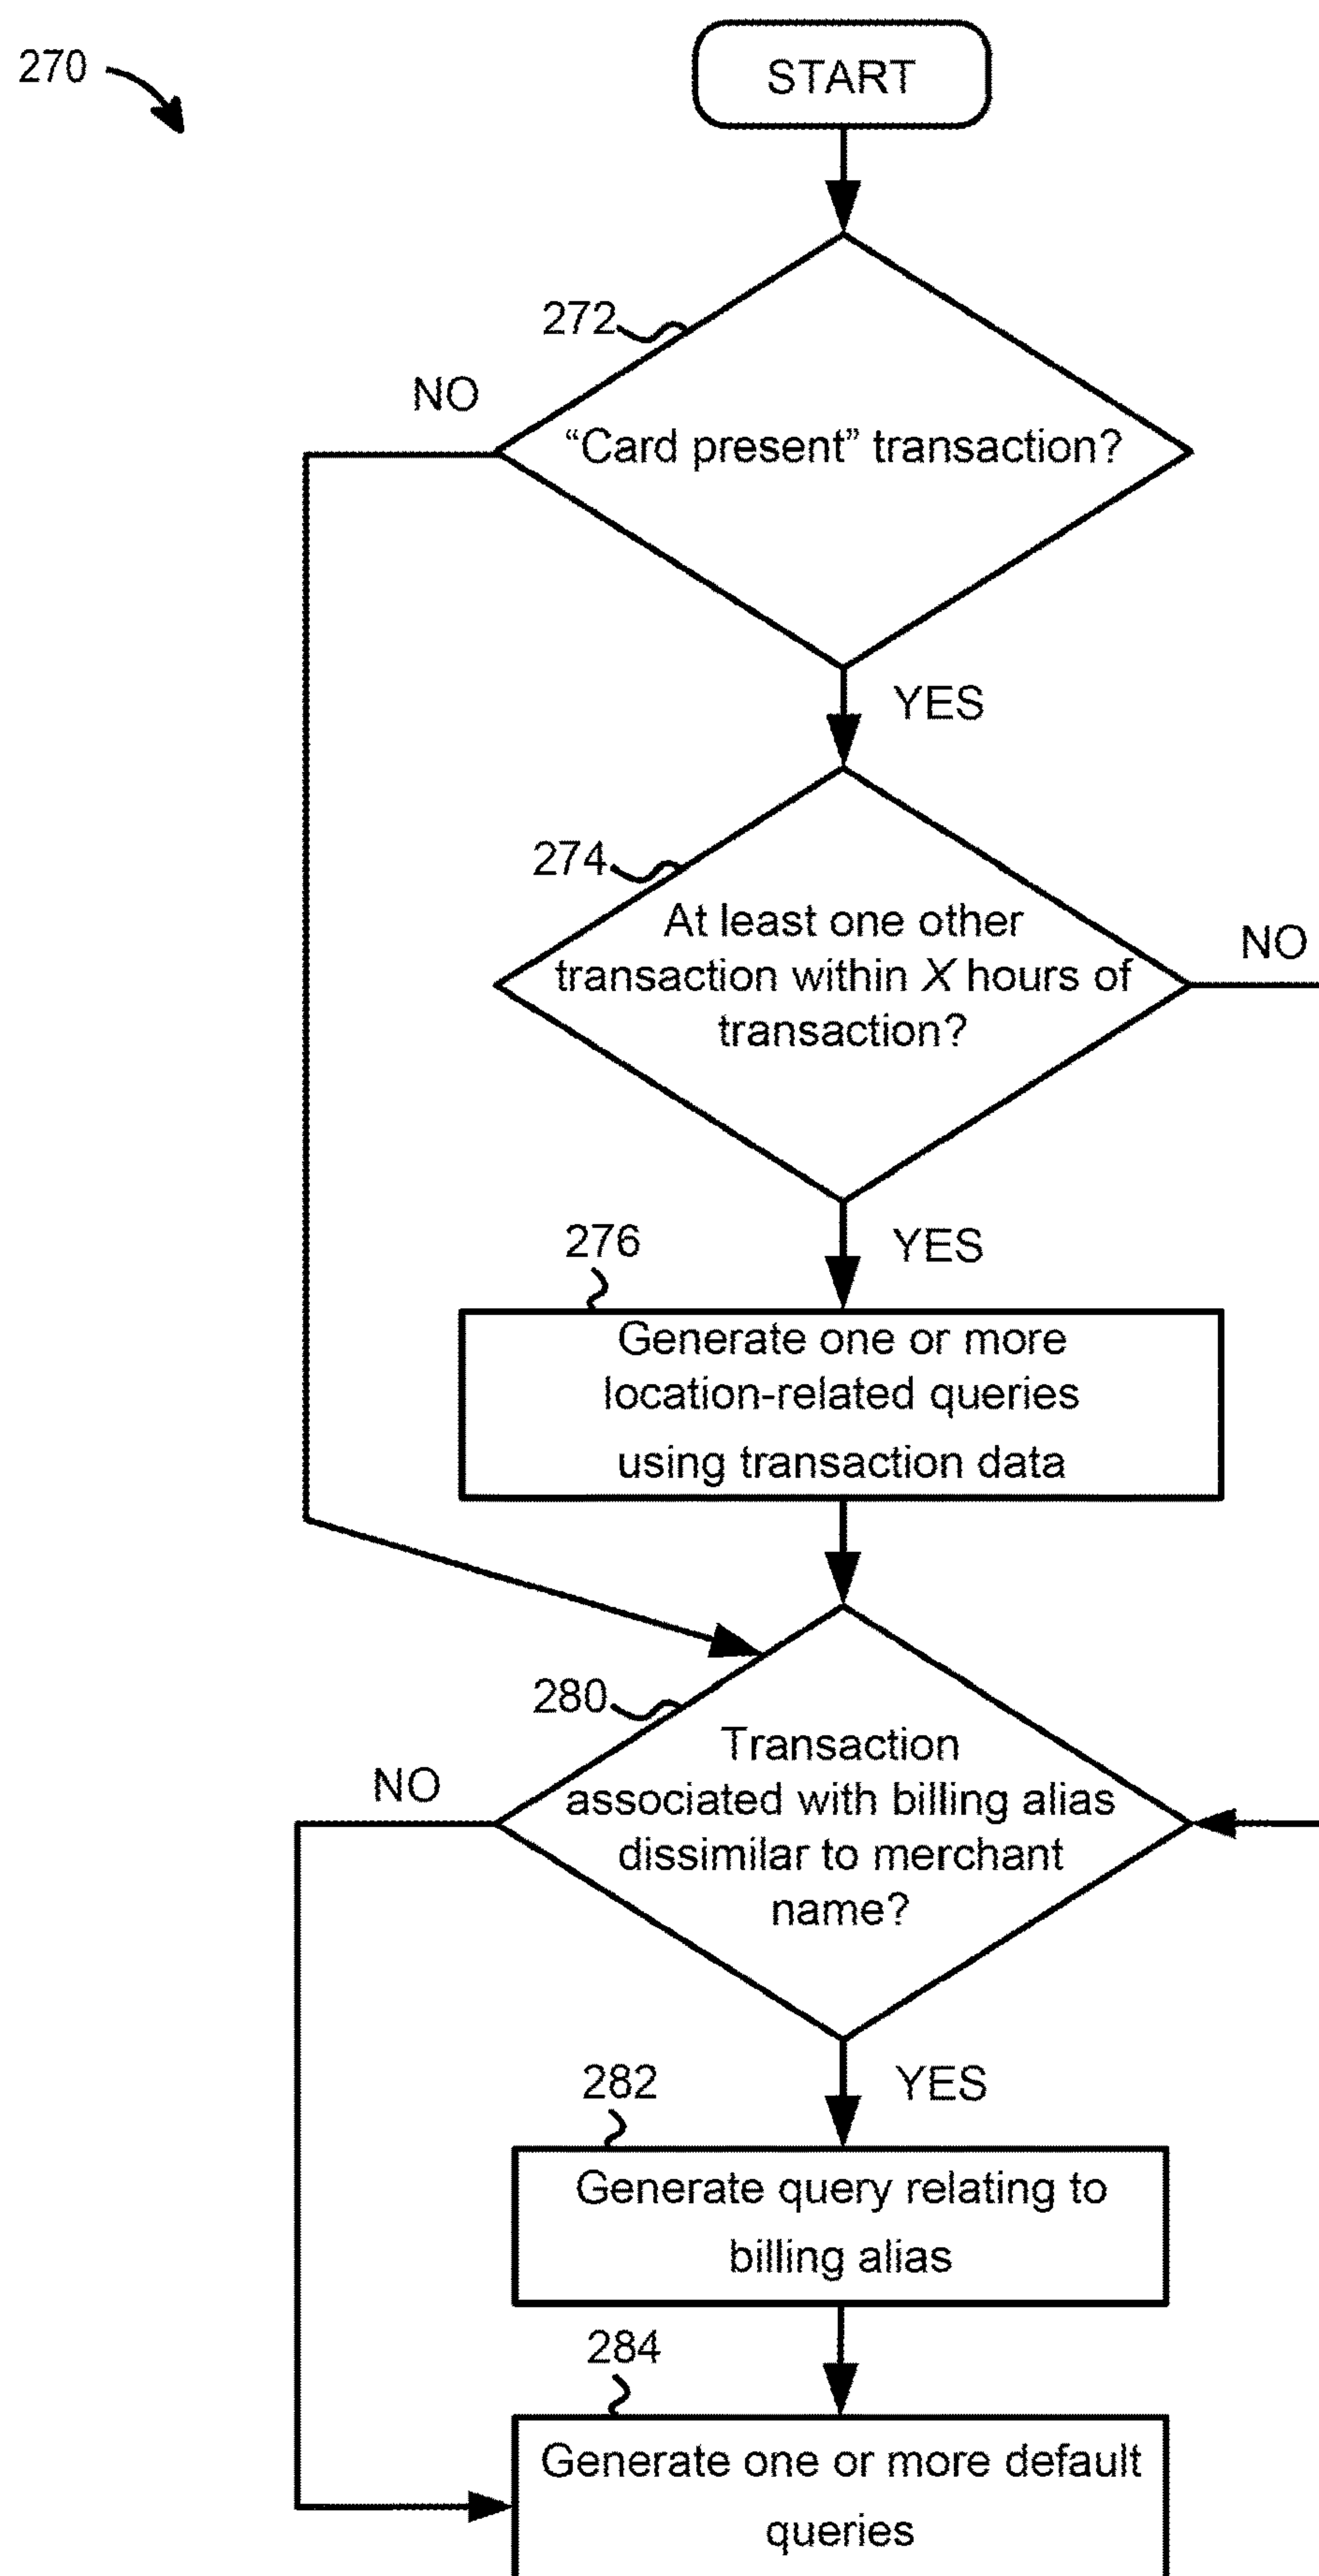
Search term also included word
“address” or “residence”

+24

Search term also included word
“employer” or “job” or “career”

+9

FIG. 4D

**FIG. 4E**


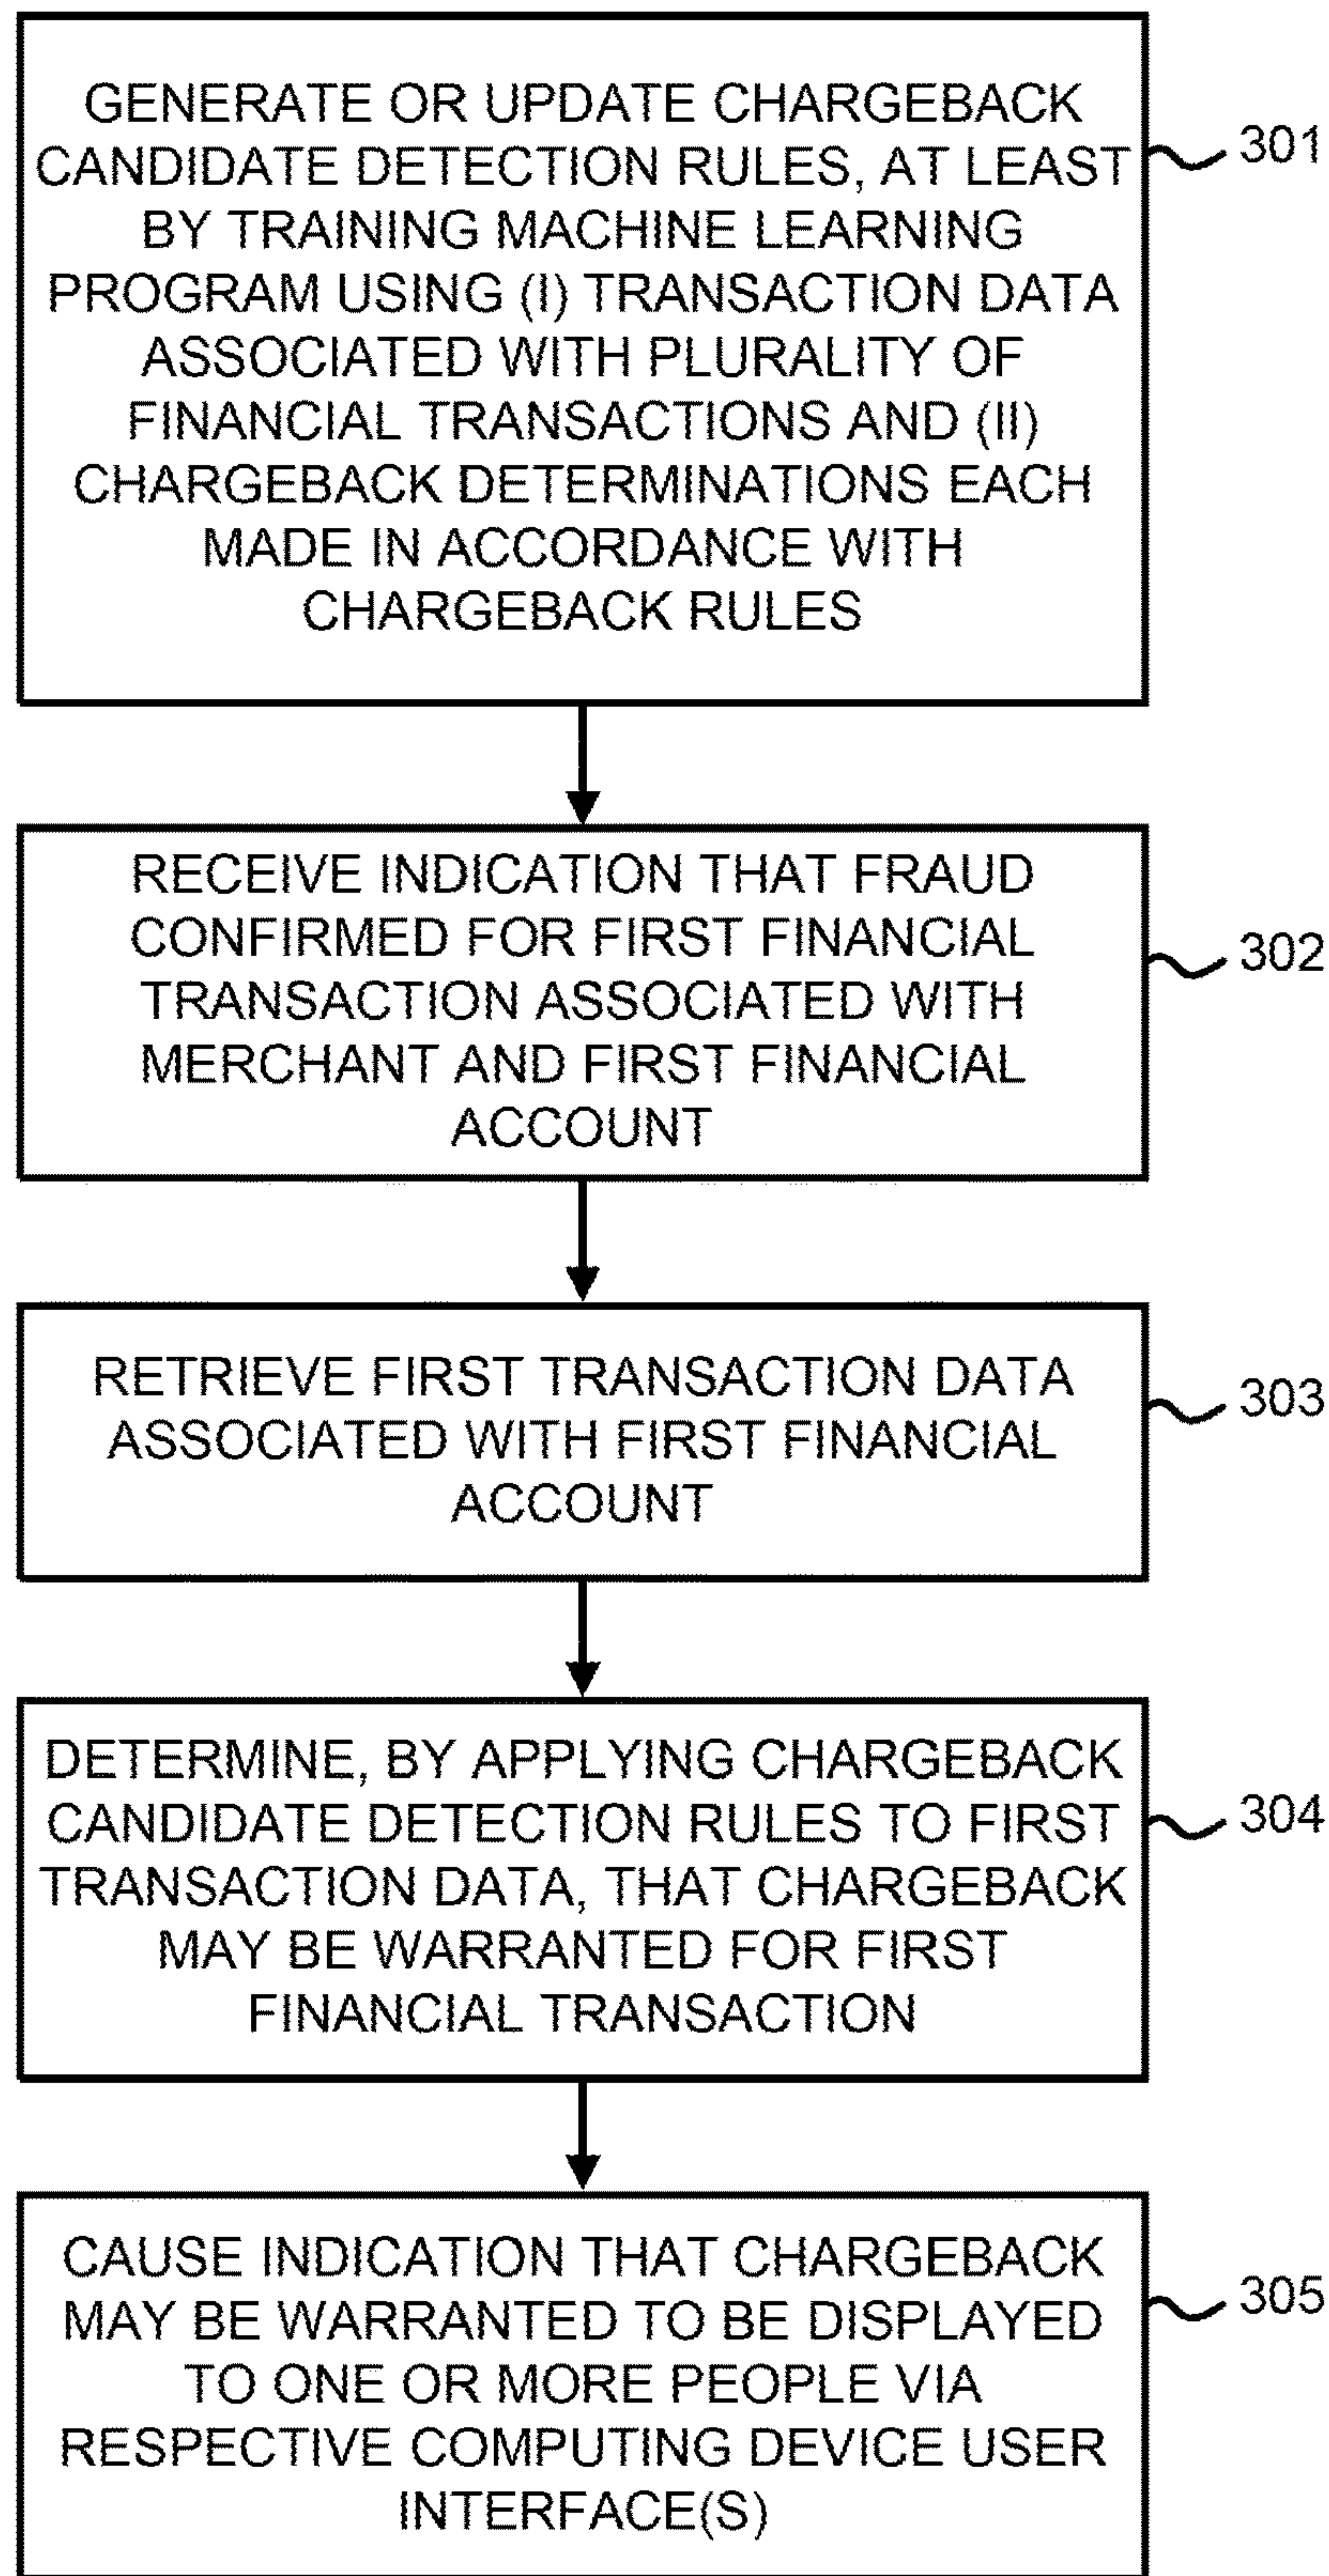

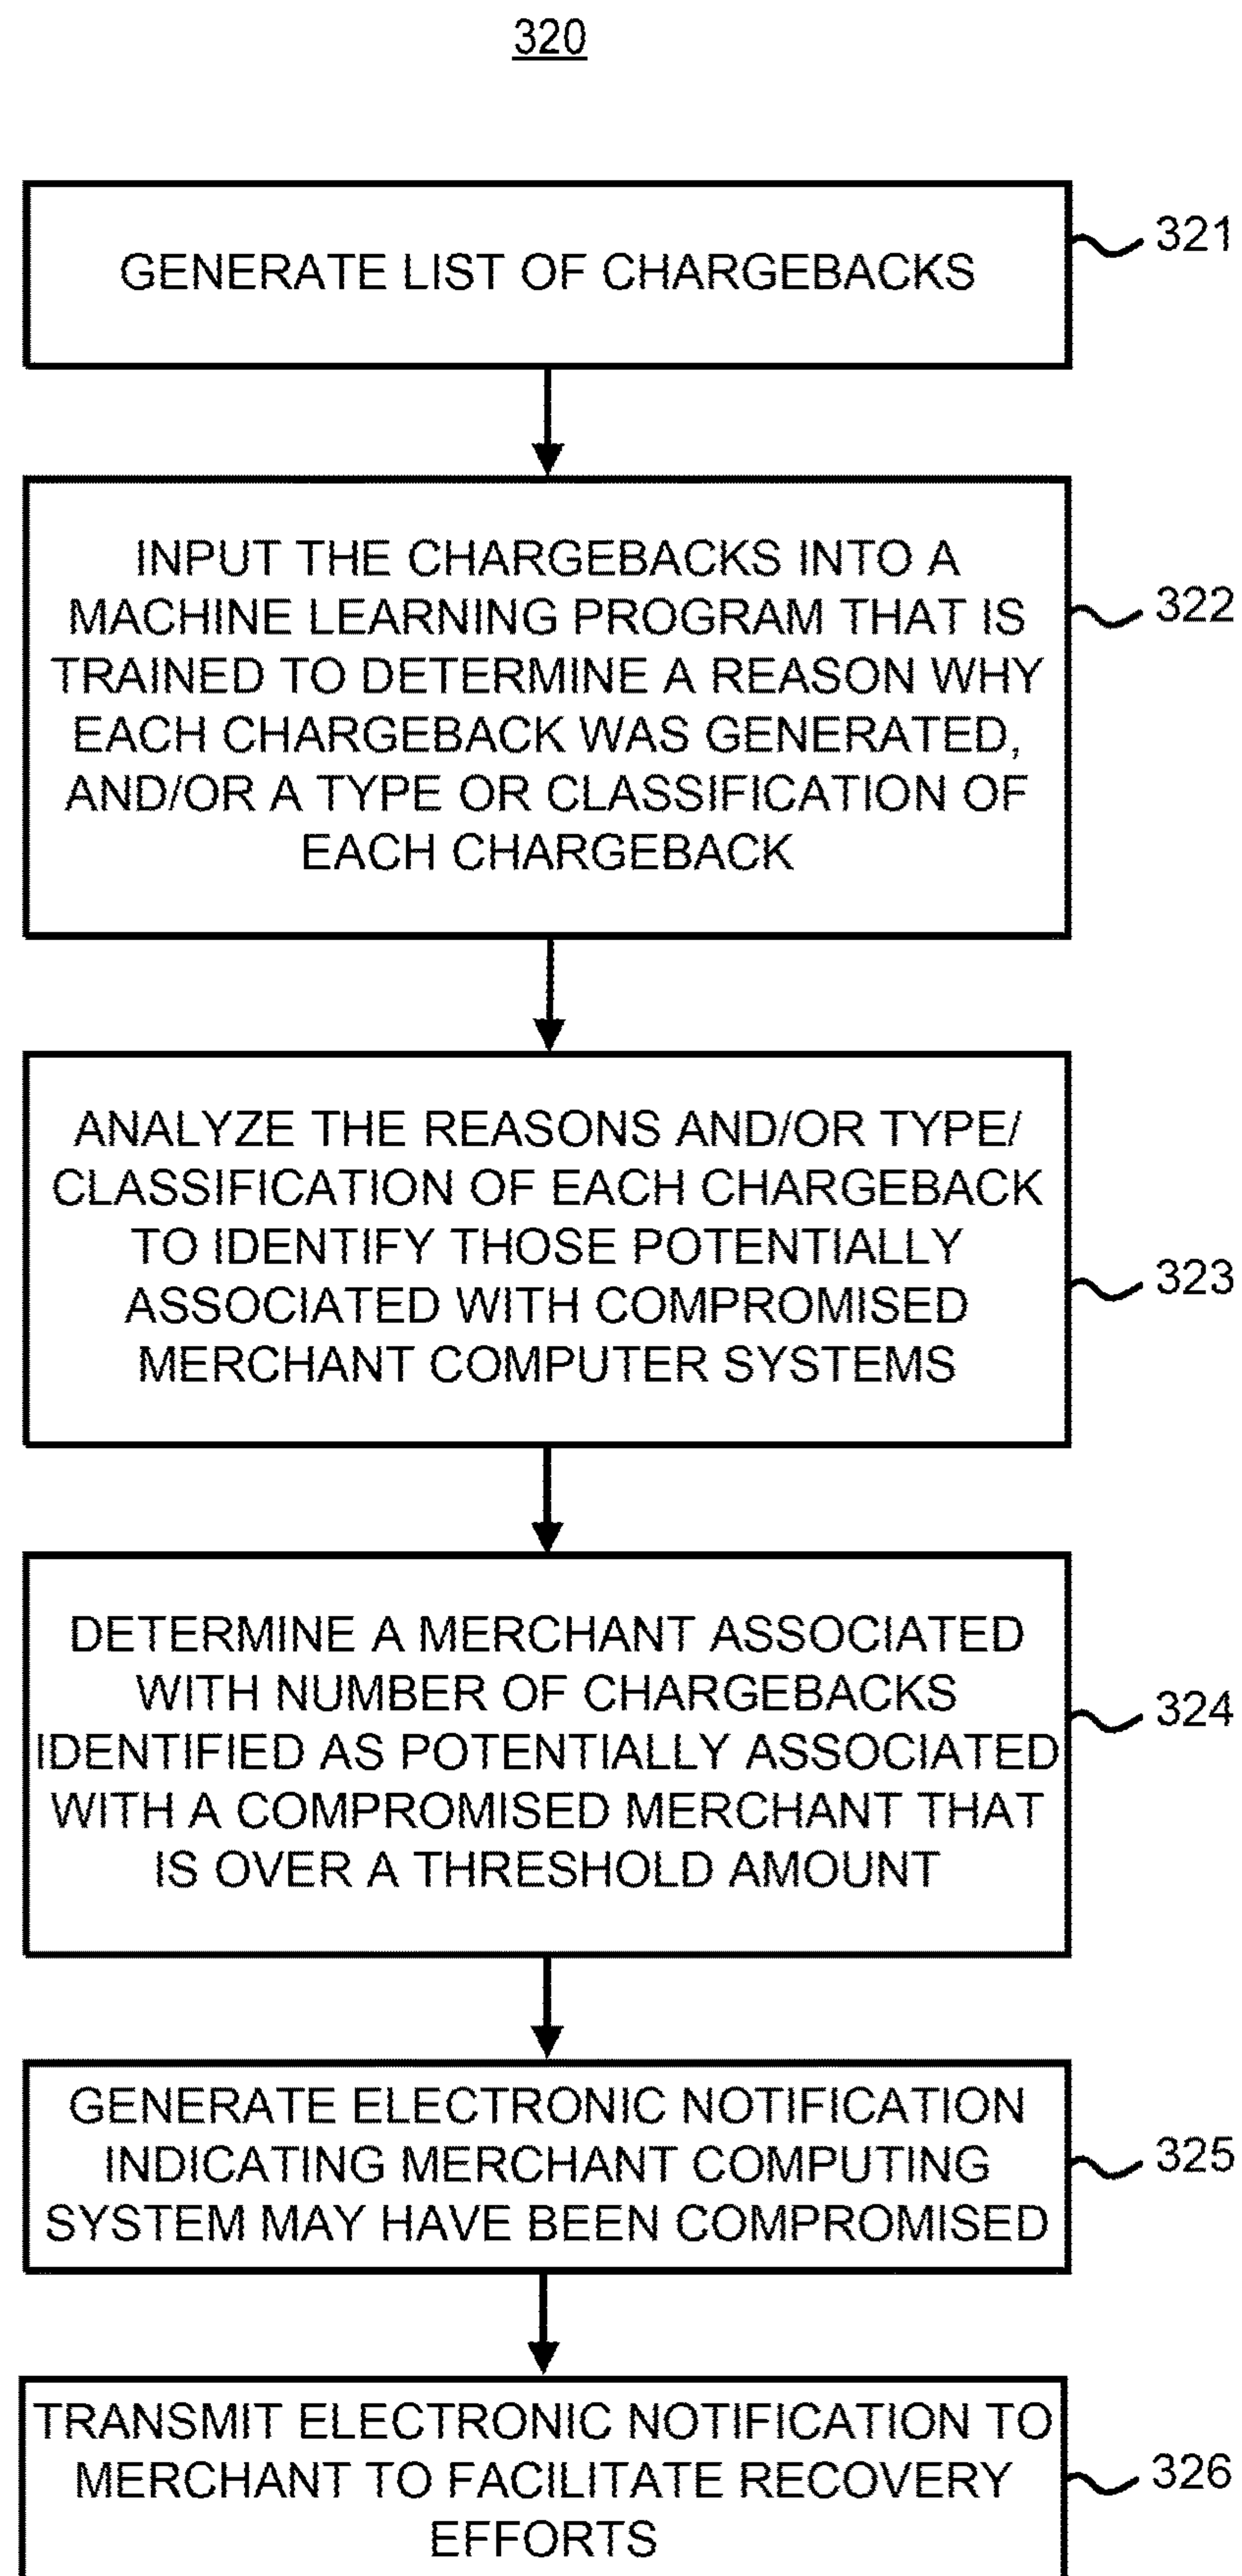
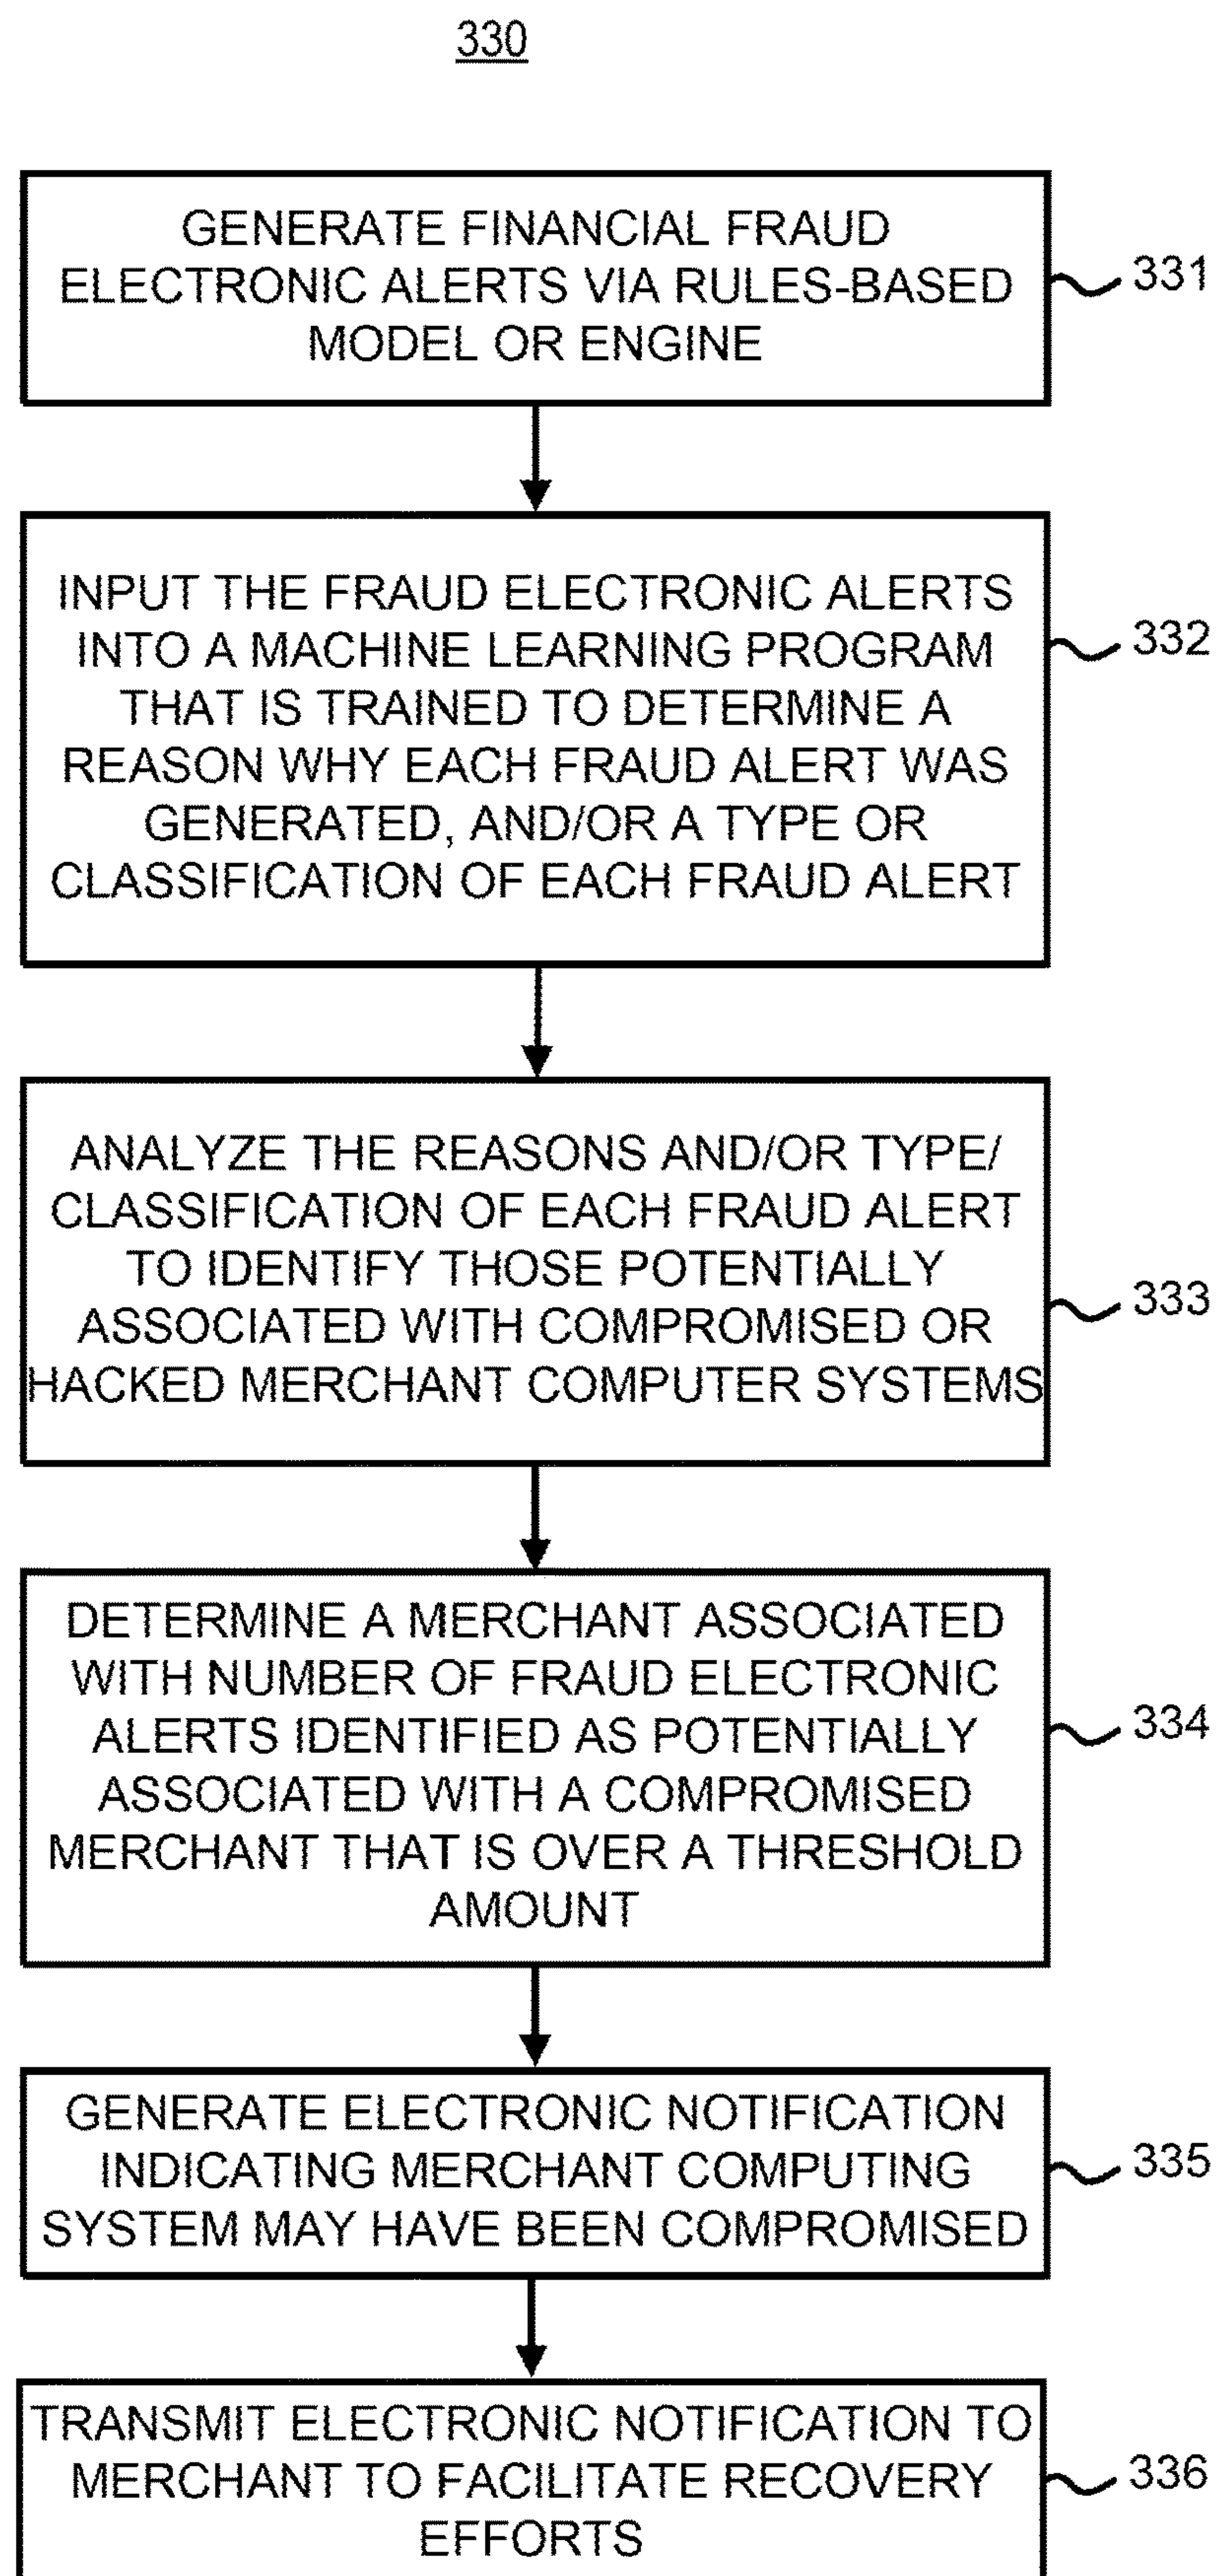


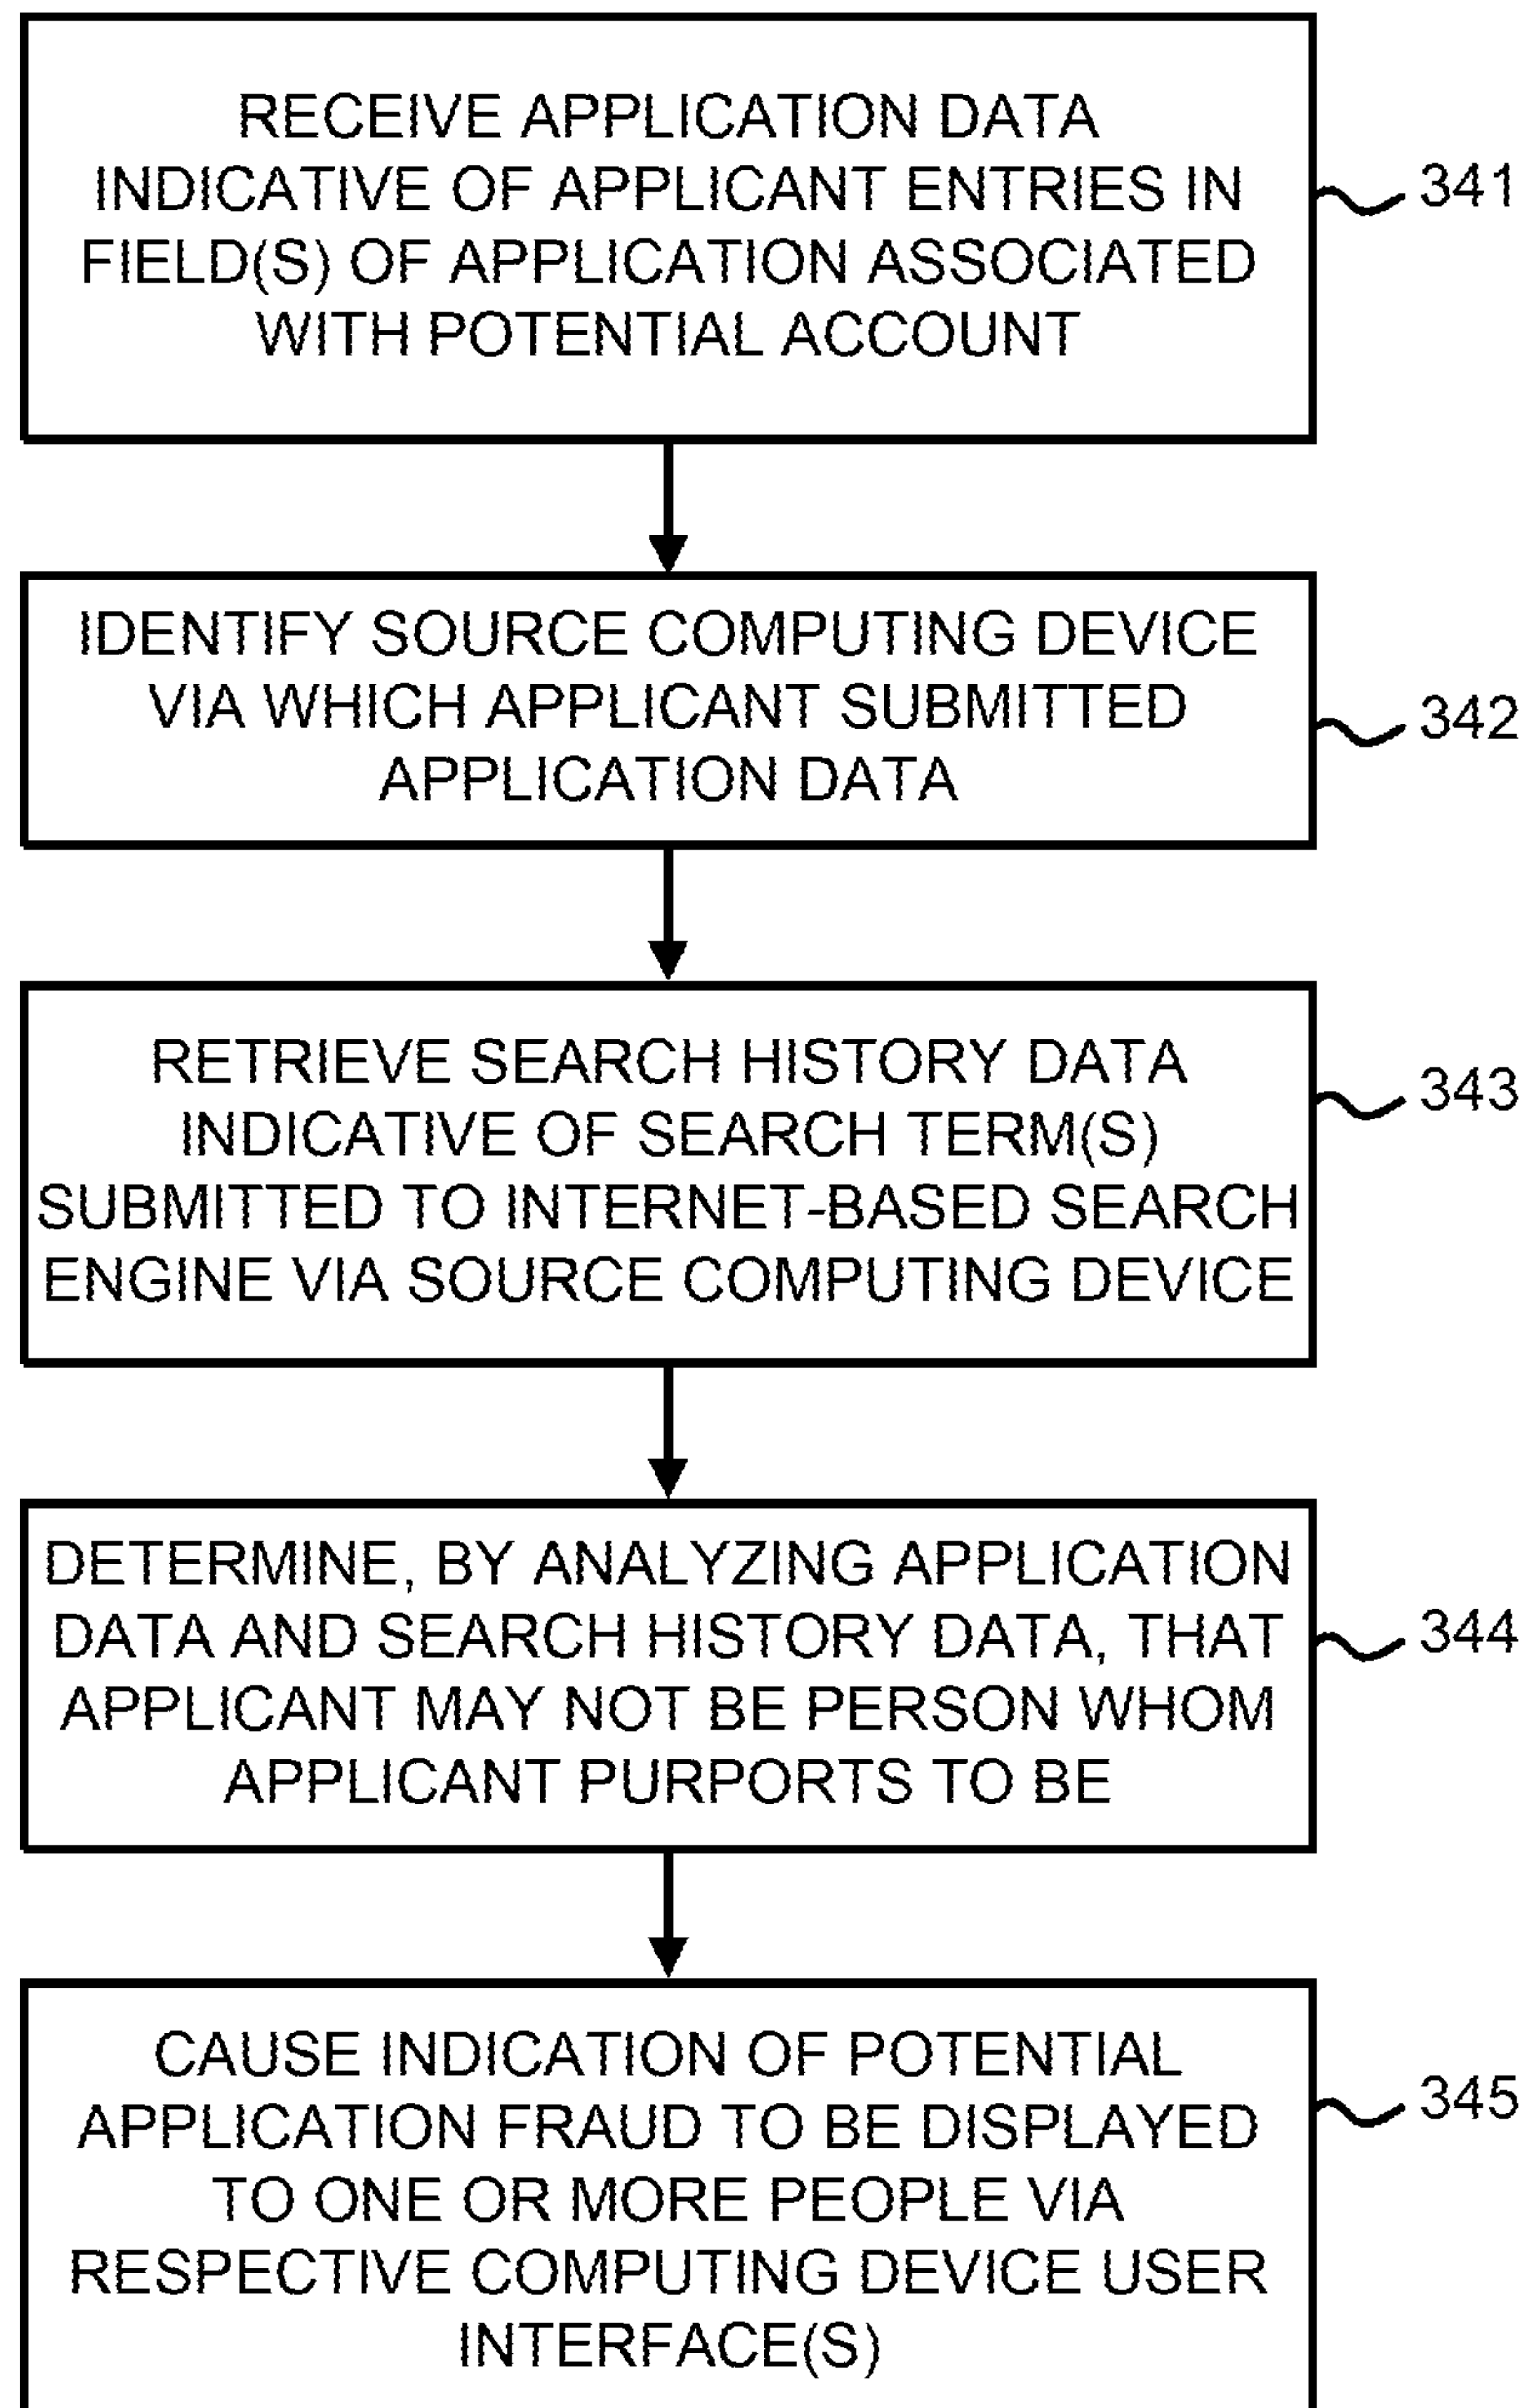
FIG. 4F

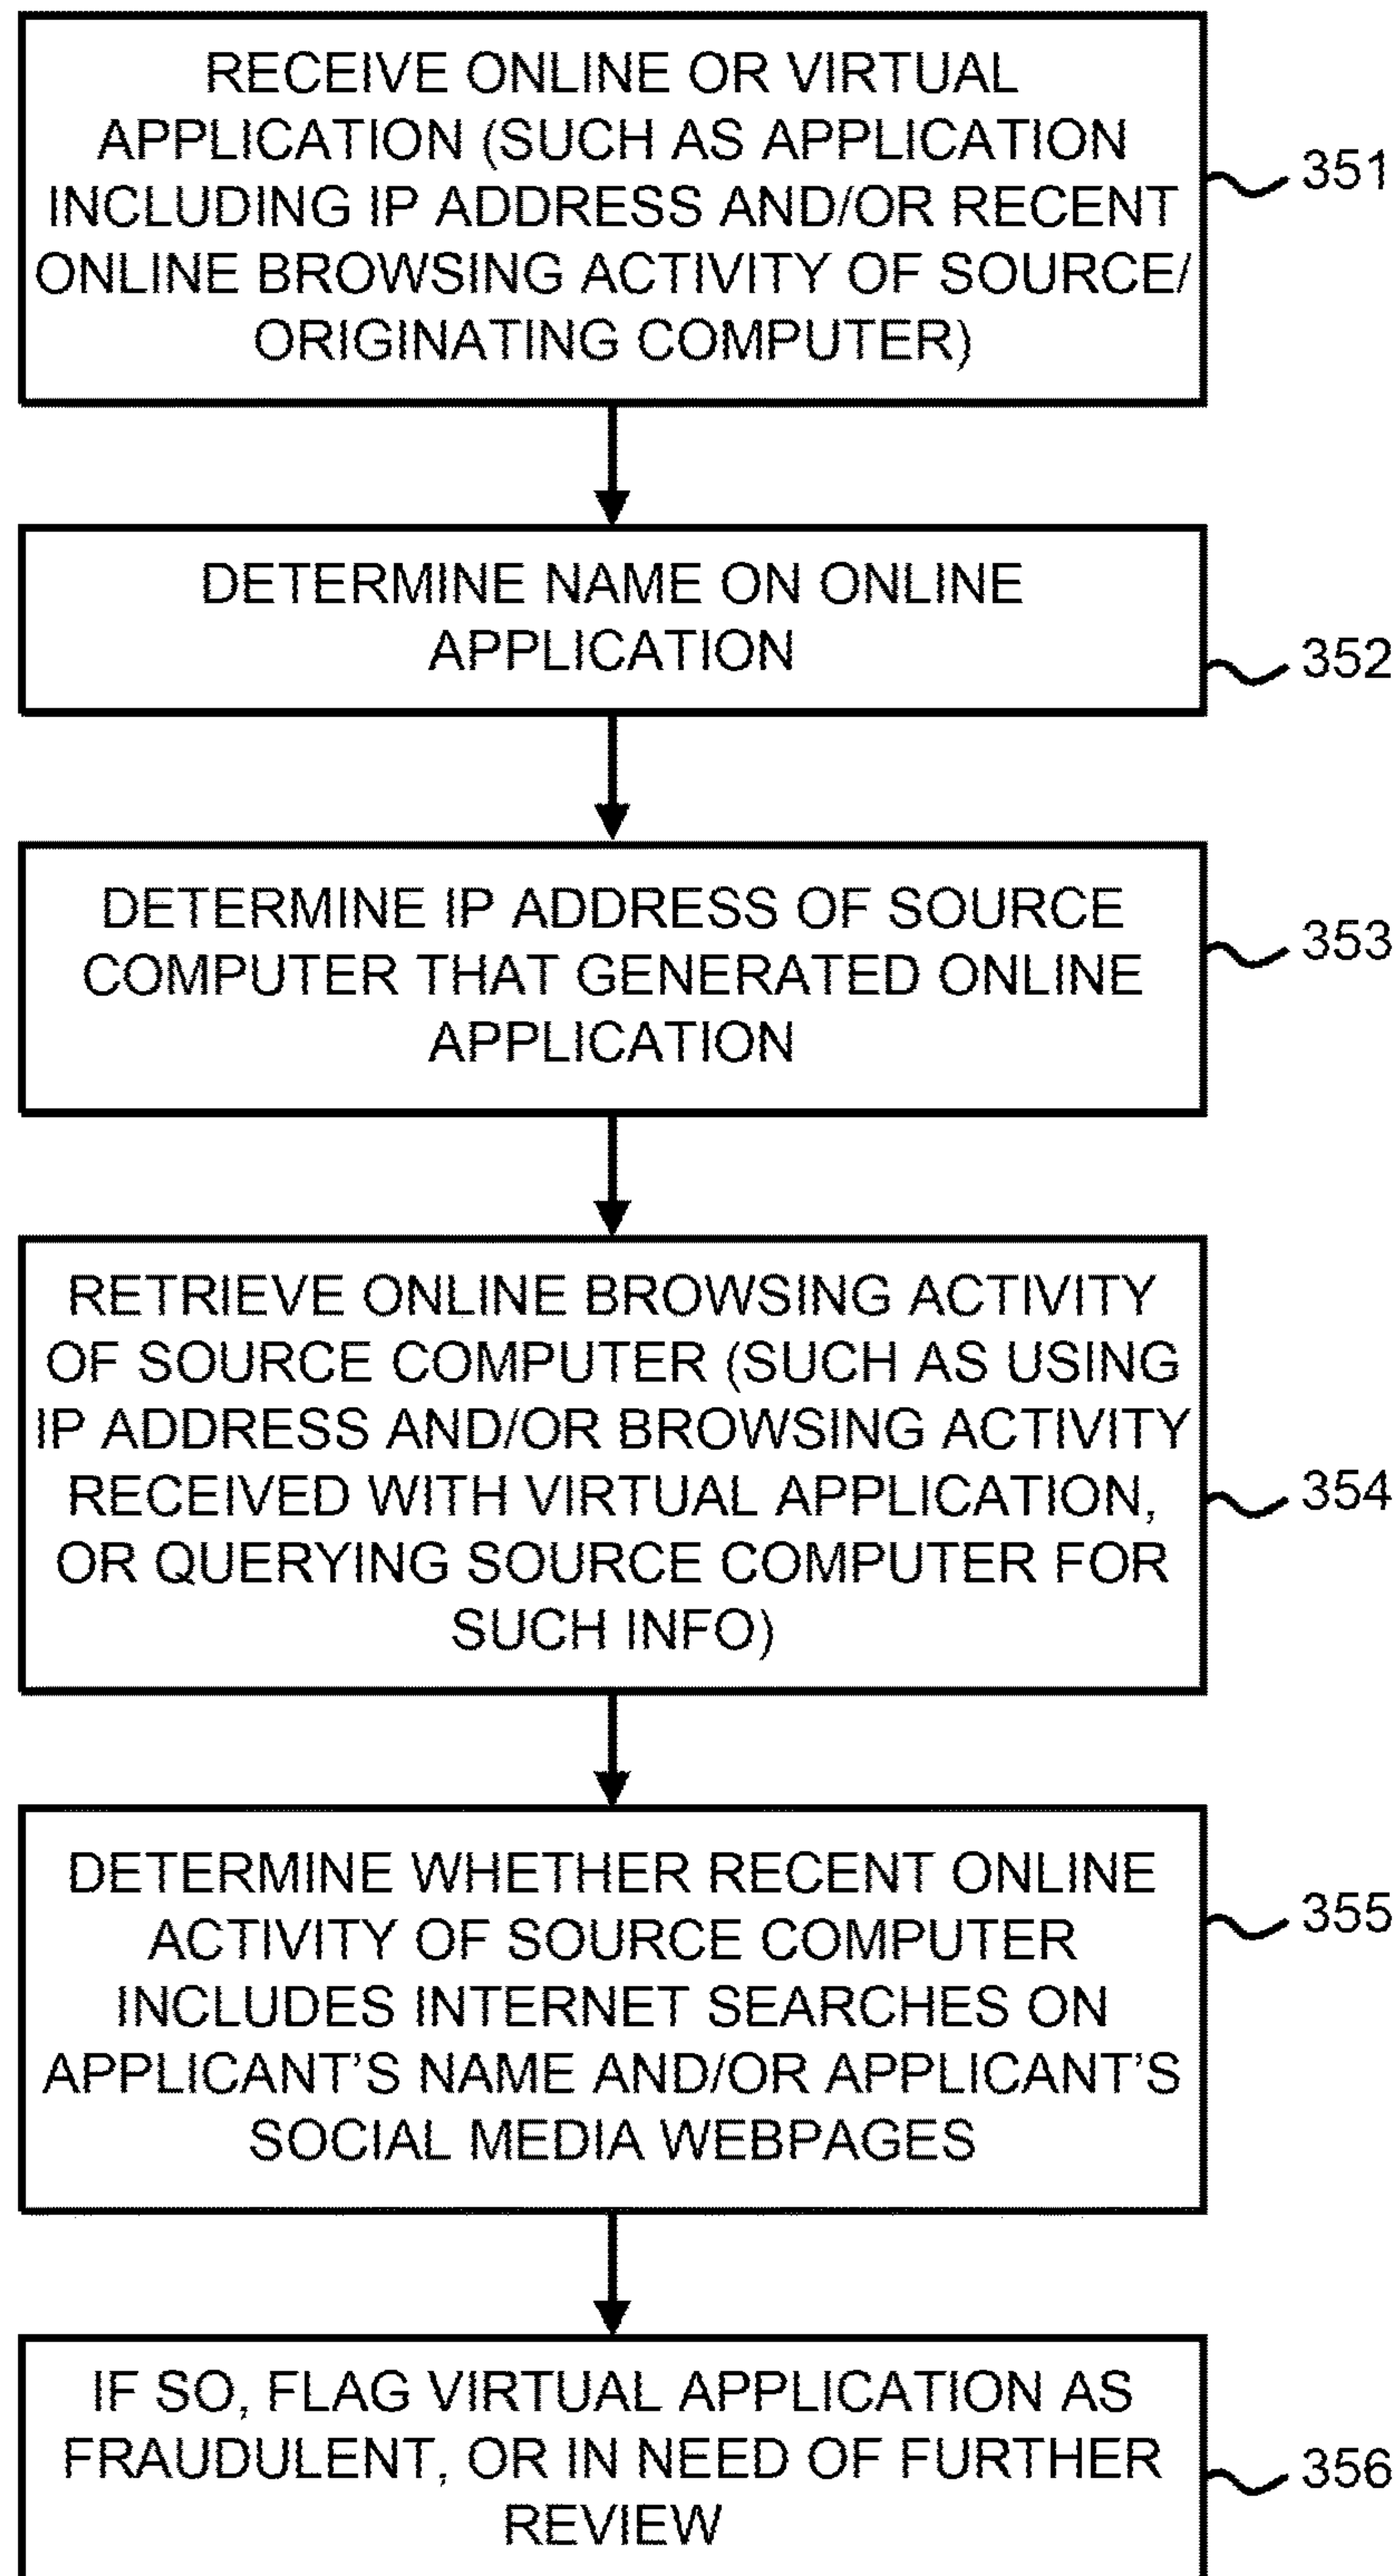
300**FIG. 5**

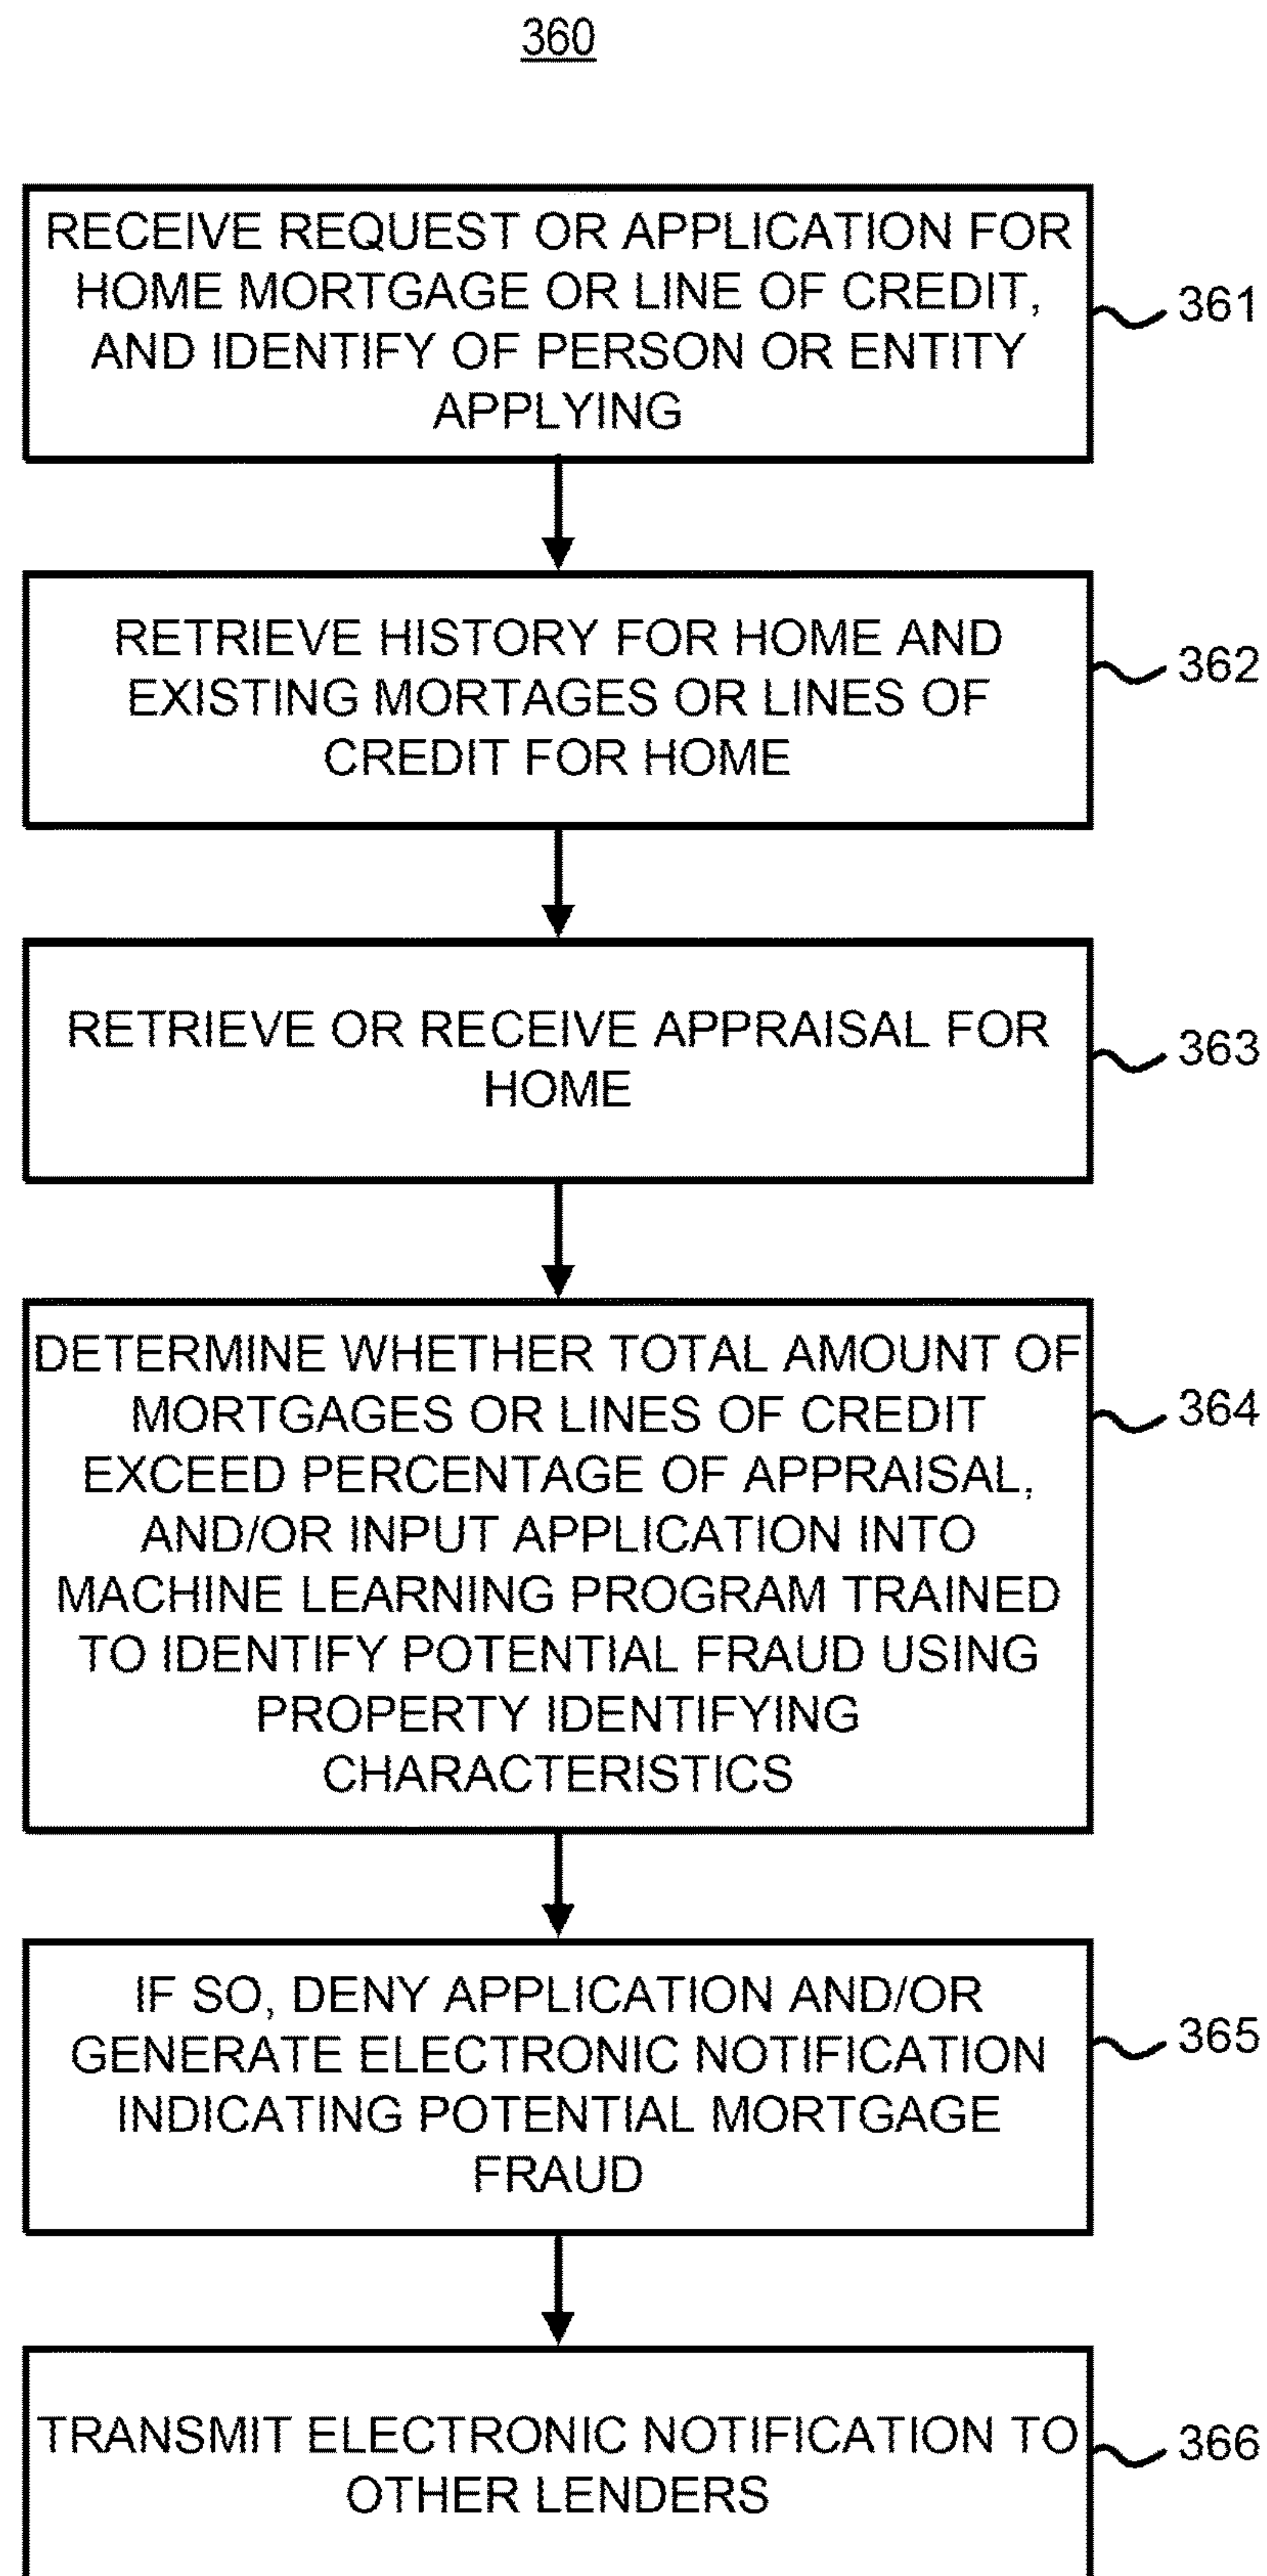
**FIG. 6**

**FIG. 7**

**FIG. 8**

340**FIG. 9**

350**FIG. 10**

**FIG. 11**




FIG. 12

IDENTIFYING FRAUDULENT ONLINE APPLICATIONS

CROSS-REFERENCE TO RELATED APPLICATIONS

This application claims the benefit of U.S. patent application Ser. No. 15/465,874, filed on Mar. 22, 2017, and entitled "Identifying Fraudulent Online Applications," which claims the benefit of U.S. Patent Application No. 62/313,196, filed on Mar. 25, 2016 and entitled "Reducing Financial Fraud Using Machine Learning and Other Techniques," U.S. Patent Application No. 62/318,423, filed on Apr. 5, 2016 and entitled "Reducing Financial Fraud Using Machine Learning and Other Techniques," U.S. Patent Application No. 62/331,530, filed on May 4, 2016 and entitled "Reducing Financial Fraud Using Machine Learning and Other Techniques," and U.S. Patent Application No. 62/365,699, filed on Jul. 22, 2016 and entitled "Detecting and/or Preventing Financial Fraud Using Geolocation Data," the disclosures of which are hereby incorporated herein by reference in their entirety.

FIELD OF THE DISCLOSURE

The present disclosure generally relates to financial fraud and, more specifically, to processing techniques that use machine learning to identify application-related to fraud (e.g., fraudulent online applications).

BACKGROUND

Financial fraud, in its many forms, is a problem of enormous magnitude and scope, causing billions of dollars in economic losses and impacting many millions of people. Types of financial fraud include use of a lost or stolen card, account takeover, skimming, chargeback ("friendly") fraud, counterfeiting, forgeries and application (e.g., loan application) fraud, to name just a few. The problem only continues to grow as various technological advances, intended to improve convenience and efficiency in the marketplace, provide new opportunities for bad actors. For example, an ever-increasing amount of fraud may be linked to online transactions made via the Internet.

Various software applications have been developed to detect potentially fraudulent transactions. For example, dollar amounts and geographic locations have generally been used to flag particular credit or debit card transactions, with cardholders then being contacted by employees of the card issuer to determine whether the transactions were indeed fraudulent. To ensure that most instances of fraud are captured, however, such techniques generally have a low threshold for triggering a fraud alert. As a result, numerous fraud alerts are false positives. The prevalence of false positives leads to a large cost in terms of the drain on human resources (e.g., calling customers to discuss each suspect transaction, and/or other manual investigation techniques), and considerable distraction or annoyance for cardholders. To provide a solution to these shortcomings in the field of automated fraud detection, innovative processing techniques capable of reducing false positives are needed.

Other conventional processes relating to financial fraud are likewise resource-intensive. For example, card issuers today perform manual reviews of fraudulent transactions to determine whether, under the lengthy and complex rules of a particular card network, chargebacks are appropriate (i.e., payments from the merchant or acquiring/merchant bank

back to the card issuer). These manual reviews are another significant drain on human resources, and may be subject to human error. Moreover, failure to properly identify chargeback scenarios may cost the card issuer money.

As another example, application fraud (e.g., obtaining a loan in the name of another person) may sometimes not be discovered until the loan has been made. By that point in time, any losses might not be fully recoverable.

BRIEF SUMMARY

The present embodiments may, inter alia, automatically detect application fraud. The detection may be based upon various types of information, and the rules used to detect such scenarios may be generated by a machine learning program.

In one embodiment, a computer-implemented method of using browsing activity to identify fraudulent online or virtual applications includes: (1) receiving, by one or more processors and via one or more transceivers, a virtual application over one or more radio frequency links; (2) determining, by the one or more processors, an applicant name on the virtual application; (3) determining, by the one or more processors, an IP address of a source computer from which the virtual application originated; (4) determining, by the one or more processors, an online browsing or search history associated with the IP address; (5) determining, by the one or more processors, whether the online browsing or search history indicates recent Internet searches for the applicant name; and/or (6) in response to determining that the online browsing or search history does indicate recent Internet searches for the applicant name, (i) flagging, by the one or more processors, the virtual application as fraudulent and (ii) generating an electronic alert indicating that the virtual application is fraudulent to facilitate identifying fraudulent virtual applications for goods or services. The method may include additional, less, or alternate actions, including those discussed elsewhere herein.

In another embodiment, a computer system is configured to use IP addresses and browsing activity to identify fraudulent virtual applications. The computer system includes one or more transceivers, and one or more processors configured to: (1) receive, via the one or more transceivers, a virtual application over one or more radio frequency links; (2) determine an applicant name on the virtual application; (3) determine an IP address of a source computer from which the virtual application originated; (4) determine an online browsing or search history associated with the IP address; (5) determine whether the online browsing or search history indicates recent Internet searches for the applicant name; and/or (6) in response to determining that the online browsing or search history does indicate recent Internet searches for the applicant name, (i) flag the virtual application as fraudulent and (ii) generate an electronic alert indicating that the virtual application is fraudulent to facilitate identifying fraudulent virtual applications for goods or services. The computer system may include additional, less, or alternate functionality, including that discussed elsewhere herein.

In another embodiment, a non-transitory, computer-readable medium storing instructions that, when executed by one or more processors, cause the one or more processors to: (1) receive, via one or more transceivers, a virtual application over one or more radio frequency links; (2) determine an applicant name on the virtual application; (3) determine an IP address of a source computer from which the virtual application originated; (4) determine an online browsing or search history associated with the IP address; (5) determine

whether the online browsing or search history indicates recent Internet searches for the applicant name; and/or (6) in response to determining that the online browsing or search history does indicate recent Internet searches for the applicant name, (i) flag the virtual application as fraudulent and (ii) generate an electronic alert indicating that the virtual application is fraudulent to facilitate identifying fraudulent virtual applications for goods or services.

BRIEF DESCRIPTION OF THE DRAWINGS

The Figures described below depict various aspects of the systems and methods disclosed herein. It should be understood that each Figure depicts an embodiment of a particular aspect of the disclosed systems and methods, and that each of the Figures is intended to accord with a possible embodiment thereof.

FIG. 1 depicts an exemplary environment in which techniques for fraud detection, verification and/or classification may be implemented, according to one embodiment.

FIG. 2 depicts an exemplary process flow for machine learning of fraud detection, verification and/or classification rules, according to one embodiment.

FIGS. 3A-3F depict exemplary process flows for machine learning of particular types of fraud detection, verification and/or classification rules, according to different embodiments.

FIGS. 4A-4F depict exemplary factors and algorithms that may be used in connection with various fraud detection, verification and/or classification rule sets, according to different embodiments.

FIG. 5 depicts a flow diagram of an exemplary method for identifying a potential chargeback scenario, according to one embodiment.

FIG. 6 illustrates an exemplary computer-implemented method of determining that a merchant computer terminal warrants a chargeback after a consumer dispute or fraudulent transaction occurs based upon dispute rules.

FIG. 7 illustrates an exemplary computer-implemented method of detecting or identifying potentially compromised merchant computing systems based upon processor analysis of electronically-generated chargebacks or actual losses.

FIG. 8 illustrates an exemplary computer-implemented method of detecting or determining potentially compromised merchant computing systems based upon processor analysis of electronically-generated fraud alerts.

FIG. 9 depicts a flow diagram of an exemplary method for detecting potential application fraud, according to one embodiment.

FIG. 10 illustrates an exemplary computer-implemented method of using IP address and/or browsing activity to identify fraudulent online or virtual applications, according to one embodiment.

FIG. 11 illustrates an exemplary computer-implemented method of detecting or identifying potential mortgage fraud.

FIG. 12 depicts an exemplary computer system in which the techniques described herein may be implemented, according to one embodiment.

DETAILED DESCRIPTION

I. Exemplary Fraud Detection and/or Classification

The embodiments described herein relate to, inter alia, wholly or partially automated detection, verification and/or classification of financial fraud. For ease of explanation, and unless otherwise clearly indicated by the context of usage, “detecting” or “determining” fraud may be used herein to

refer to initially flagging fraudulent (or potentially fraudulent) activity, to verifying/confirming that suspect/flagged activity was indeed fraudulent, or generally to both. The systems and techniques described herein may be used, for example, to identify, prevent and/or quantify/measure instances of lost or stolen card use, account takeover, counterfeiting, skimming, chargeback (“friendly”) fraud, collusive merchant fraud, application (e.g., loan application) fraud, mortgage fraud, and/or one or more other types of fraud relating to existing and/or potential financial transactions and/or accounts. Moreover, those skilled in the art will appreciate that at least some of the technical advancements described below (and/or shown in the accompanying figures) are not necessarily restricted to the financial field.

In some embodiments, a fraud detection and/or classification system may analyze data relating to a number of existing or potential financial accounts. The analysis/processing may be performed in batch processing operations, or substantially in real-time (e.g., as the data is generated and/or as financial transactions occur, etc.), and the data may be obtained from a variety of sources based upon the particular embodiment and/or scenario. In one embodiment, for example, data from financial account records may be analyzed, along with data indicating online activity of an account holder, location data (e.g., global positioning satellite (GPS) data from a smartphone or vehicle of the account holder) and/or other data, to determine whether a particular financial transaction was fraudulent or likely fraudulent. The analysis may be performed automatically after the transaction has been made, or may be performed in response to a person or algorithm flagging the transaction as a potentially fraudulent one, for example.

The analysis may include determining whether the account holder has expressed interest in the object (e.g., product or service) of the transaction or the merchant, and/or determining whether the transaction is consistent with spending patterns associated with the account holder (e.g., spending patterns identified using the account holder’s transaction records), for example. In the case of multiple account holders (e.g. multiple credit or debit card holders), accuracy may be improved by identifying spending patterns at the individual level rather than, or in addition to, at the aggregate account level. For example, a maximum amount of money typically spent in a single transaction (e.g., over the course of a one-month window, etc.) may be determined for each of two cardholders listed on a single account, and the maximum amount for the cardholder who purportedly made a particular purchase may be compared to the purchase amount to determine whether fraud is suspected.

In another exemplary embodiment, financial transaction data may be analyzed to determine whether a chargeback payment from the merchant or acquiring bank to a card issuer may be appropriate in connection with a particular fraudulent transaction. For example, the card information entry mode (e.g., collecting card information by inserting the card in a chip reader, swiping the card, manually entering the card information, etc.), the transaction amount, the similarity to other transaction(s), and/or other information may be used to identify which fraudulent transactions are relatively strong chargeback candidates. The analysis may be performed in response to a cardholder reporting the transaction as fraudulent, or after a card issuer has confirmed that the transaction was fraudulent, for example. For the subset of instances where a fraudulent transaction has been identified as a chargeback candidate, a full set of chargeback rules (e.g., devised by a card network entity such as VISA®, Mastercard®, American Express®, Discover®, etc.) may be

5

manually or automatically applied to determine whether a chargeback process should be initiated (or continued).

In another exemplary embodiment, application data (e.g., information entered in fields of an online application) may be analyzed in conjunction with search terms entered by a user at a computing device (e.g., the device from which the user submitted the application information) to determine whether the person proffering the application is not the person that he or she purports to be. For example, if the person submitting an application had previously used an Internet-based search engine to search for results associated with the purported applicant's name (e.g., by using the name as a search term, possibly in addition to other terms such as "address" and/or "employer," etc.), the application may be flagged for suspected fraud, and subjected to additional steps of manual and/or automated review.

In another exemplary embodiment, a fraud dispute resolution process (e.g., after a customer has reported a fraudulent or unrecognized transaction associated with his or her account) may be facilitated using machine learning techniques. For example, a machine learning program may be trained, using past dispute resolution interactions with customers and the associated outcomes (fraud determinations), to identify various types of information that, if elicited from customers, tend to be indicative of fraud or the absence thereof. When fraud is suspected for a particular transaction, one or more queries for the individual purportedly making the transaction may be automatically generated using the types of information identified by the machine learning program, as well as information about the suspect transaction and/or related transactions (e.g., dates, locations, amounts, etc.). In some embodiments and/or scenarios, responses to the queries may be collected and analyzed to automatically generate additional queries, with the end goal of discerning whether the transaction was authorized. For example, queries may include asking whether a cardholder recalls particular other transactions that appear on the cardholder's account and were made around the same time as the suspect transaction (and/or from the same merchant), asking whether the cardholder recalls being in a particular location at a particular time (e.g., a location associated with another transaction appearing on the cardholder's account), whether the cardholder is aware of a particular billing alias used by a merchant, and so on.

In another exemplary embodiment, image data corresponding to a particular physical document (e.g., a personal or cashier's check, a driver's license or other identification card, etc.) may be analyzed, using rules generated by a machine learning program, to determine whether the document is, or may be, fraudulent (e.g., a counterfeit document, and/or a document that includes forged contents). For example, the machine learning program may be trained using images of multiple other documents, and fraud determinations made in connection with those other documents. The machine learning program may learn which ranges and/or tolerances for dimensions, fonts, colors, patterns, etc., tend to be most indicative of counterfeiting, for example. A forgery may be detected based upon factors relating to the contents of various fields in a document, such as whether handwriting, a signature, and/or a date format (e.g., "Jan. 1, 2016," "1/1/16," etc.) matches that used for other personal checks from a particular account holder, for example. The fraud determination may be made substantially in real-time to provide a warning, if needed, to a merchant making a sale, for example, or may be used to flag a relatively small number of documents for physical review at a later time, etc.

6

In another exemplary embodiment, machine learning techniques may be used to analyze financial transactions for purposes of classifying potentially fraudulent behavior (e.g., "counterfeiting," "skimming," "lost or stolen card," etc.).

For example, the machine learning program may be trained using fraud classifications made in connection with multiple other financial accounts. The machine learning program may learn which types of data tend to be indicative of different classifications (e.g., transaction amount, credit card information entry mode, particular types of online activity data, etc.), and/or which data values tend to be indicative of different classifications (e.g., transactions over \$10,000, manual card number entry, etc.), for example. Once a class of potential fraud has been identified for a particular transaction, the classification may be used to facilitate or guide a further, more in-depth analysis or investigation. Alternatively, or in addition, the classification may be used to calculate one or more metrics indicating the prevalence of that type of fraud.

By replacing conventional processing techniques with one or more of the processing techniques described herein, problems that have beset the field of fraud detection, classification and/or prevention in the past may be greatly mitigated or eliminated. For example, information that has conventionally been overlooked or ignored may be used to more accurately detect, prevent and/or classify fraud, and/or to reduce false positive fraud alerts. As another example, a significant amount of time may be saved by removing the need for manual investigations, or by reducing the number of instances where manual investigations are required.

II. Exemplary Environment for Implementing Fraud Detection and/or Classification Processing Techniques

FIG. 1 depicts an exemplary environment 10 in which techniques for fraud detection and/or classification may be implemented, according to one embodiment. The environment 10 may include an anti-fraud services system (AFSS) 12, a financial account management system (FAMS) 14, a card network computing system 16, a number of cardholder computing devices 20, a number of merchant computing systems 22, a number of other sources 24, and a network 26. It is noted that, in other embodiments and/or scenarios, the environment 10 may include more, fewer and/or different components than those shown in FIG. 1, such as any of those discussed elsewhere herein. For example, the environment 10 may include one or more additional financial account management systems and/or card network computing systems, and/or one or more of the cardholder computing devices 20 may instead be a computing device of a holder of a non-card account (e.g., a checking, savings or loan account) or an applicant for a new account (e.g., a new loan account). As another example, the environment 10 may include a computing system of one or more acquiring/merchant banks, and some or all of the communications with merchant computing systems 22 described below may instead be with the acquiring bank(s).

FAMS 14 may be associated with (e.g., owned and/or maintained by) a bank or other financial entity. For example, FAMS 14 may be a bank that acts as a card issuer associated with a particular type of card network (e.g., VISA®, Mastercard®, etc.), and/or an entity that provides loans (e.g., mortgage, home equity, vehicle, etc.), saving/checking account services, and/or other financial services to customers. FAMS 14 may maintain an account records database 30 that stores various kinds of account information, including account holder information (e.g., names, addresses, etc.) and data indicative of financial transactions made in connection with each account (e.g., dates, amounts and merchants for

credit or debit card transactions, dates and amounts for customer deposits and withdrawals, etc.). Account records database 30 may store account information for some or all of the cardholders associated with cardholder computing devices 20, for example. While shown in FIG. 1 as a single entity within FAMS 14, it is understood that account records database 30 may, in some embodiments, be distributed across multiple databases and/or multiple physical/hardware memories, and/or may be wholly or partially external to (e.g., remote from) FAMS 14.

AFSS 12 may generally provide services that help to detect and/or classify fraudulent activity in connection with existing and/or potential (e.g., applied for) financial accounts, such as the accounts managed by FAMS 14. In some embodiments, AFSS 12 is included within FAMS 14. As seen in FIG. 1, AFSS 12 may include a network interface 32, a memory 34, and a fraud detection/classification unit 36.

Network interface 32 may include hardware, firmware and/or software configured to enable AFSS 12 to wirelessly exchange electronic data with one or more other components of environment 10 via network 26. For example, network interface 32 may include an Ethernet port, a modem, a router, and/or one or more other ports and/or transceivers for one or more other wired and/or wireless communication technologies.

Memory 34 may be a computer-readable, non-transitory storage unit or device, or collection of units/devices, and may include persistent (e.g., hard disk) and/or non-persistent memory components. Memory 34 may store instructions that are executable on one or more processors of AFSS 12 (not shown in FIG. 1) to perform various operations, including the instructions of various software applications and data generated and/or used by such applications.

Card network computing system 16 may be a computing system (e.g., one or more servers) of a credit and/or debit card network entity, such as VISA® or Mastercard®, for example. In some embodiments and/or scenarios where the card network entity also acts as the issuer (e.g., American Express® or Discover®), card network computing system 16 may include FAMS 14. Card network computing system 16 may provide various services to FAMS 14 and/or AFSS 12. For example, card network computing system 16 may provide electronic updates to chargeback rules, fraud scores for particular customers and/or transactions, and so on.

Each of cardholder computing devices 20 may be a computing device of a respective holder of a credit or debit card account managed by FAMS 14. For example, one or more of cardholder computing devices 20 may be desktop computers, laptop computers, tablet computers, smartphones, smart watches, and so on. The cardholders (e.g., credit or debit card account holders) may use cardholder computing devices 20 to access (e.g., view, modify, etc.) their account information stored in account records database 30 online via network 26. In some embodiments where AFSS 12 detects and/or classifies activity not related to credit or debit card fraud (e.g., a fraudulent application for a home equity loan, etc.), cardholder computing devices 20 may instead be computing devices of other types of customers or potential customers, such as holders of non-card-based accounts, or individuals who have submitted an online application for a loan, etc., as discussed further below. In some of these embodiments, the environment 10 may omit card network computing system 16.

Each of merchant computing systems 22 may include one or more computing devices associated with a particular provider of products and/or services. For example, some or

all of merchant computing systems 22 may include servers associated with online retailers. Alternatively, or additionally, some or all of merchant computing systems 22 may include point-of-sale terminal devices providing credit and/or debit card payment processing features for “card present” transactions. In some embodiments where AFSS 12 detects and/or classifies activity not related to customer purchases (e.g., if AFSS 12 only detects loan application fraud, etc.), the environment 10 may omit merchant computing systems 22.

The other sources 24 may include computing devices and/or systems associated with sources of one or more other types of information. For example, other sources 24 may include vehicle telematics systems (e.g., installed in vehicles of cardholders associated with cardholder computing devices 20), one or more Internet service providers (ISPs) (e.g., ISPs providing Internet access to some or all cardholders), “smart home” system devices (e.g., installed in homes of some or all cardholders), and/or other systems/devices. In some embodiments, the environment 10 does not include the other sources 24.

Network 26 may communicatively couple some or all of the components shown in FIG. 1. For example, FAMS 14 may use network 26 to communicate with AFSS 12, card network computing system 16, cardholder computing devices 20 and/or merchant computing systems 22. As another example, AFSS 12 may use network 26 to communicate with FAMS 14, card network computing system 16, cardholder computing devices 20, merchant computing systems 22 and/or one or more of the other sources 24. While shown as a single entity in FIG. 1, network 26 may include multiple communication networks of one or more types (e.g., one or more wired and/or wireless local area networks (LANs), and/or one or more wired and/or wireless wide area networks (WANs) such as the Internet). Moreover, network 26 may use partially or entirely distinct network components to support communications between different endpoints or computing devices, such as wireless communication or data transmission over one or more radio frequency links and/or wireless communication channels. For example, the portion(s) of network 26 used for communications between FAMS 14 and AFSS 12 may be the same as, or different than, the portion(s) of network 26 used for communications between FAMS 14 and one or more of cardholder computing devices 20 over one or more radio links or wireless communication channels, or between AFSS 12 and one or more of the other sources 24, etc. Those skilled in the art will appreciate different types of networks that are appropriate for network 26, depending upon, for example, how AFSS 12, FAMS 14 and/or other components of environment 10 are localized or distributed across a relatively large geographic area.

Generally, fraud detection/classification unit 36 of AFSS 12 may detect fraudulent activity, confirm whether suspected or reported fraudulent activity is truly fraudulent, and/or classify fraudulent or suspected fraudulent activity. For example, fraud detection/classification unit 36 may analyze each transaction stored in account records database 30 to determine whether that transaction is, or potentially is, fraudulent. Alternatively, fraud detection/classification unit 36 may analyze only those transactions that were flagged as possibly being fraudulent (e.g., by a cardholder calling in to report an unauthorized and/or unrecognized transaction, or by FAMS 14 or AFSS 12 generating a preliminary fraud alert after applying an initial set of rules to a transaction, etc.). Fraud detection/classification unit 36 may also, or instead, support additional functionality, such as that

described below in connection with the various components of fraud detection/classification unit 36 shown in FIG. 1.

As seen in FIG. 1, fraud detection/classification unit 36 may include a machine learning (ML) rule generator 40, an external data collection unit 42, a behavior analysis unit 44, a dispute resolution unit 46, a chargeback analysis unit 50, an image analysis unit 52, a classification unit 54, and/or a notification unit 56. In other embodiments, fraud detection/classification unit 36 may include more, fewer and/or different components/units than those shown in FIG. 1. In some embodiments, each of ML rule generator 40, external data collection unit 42, behavior analysis unit 44, dispute resolution unit 46, chargeback analysis unit 50, image analysis unit 52, classification unit 54, notification unit 56, and/or other units or components of fraud detection/classification unit 36 may be a software component stored in memory 34 and implemented by one or more processors of one or more computing devices (e.g., servers) included in AFSS 12.

ML rule generator 40 may generally analyze various types of data to generate and/or update fraud detection and/or classification rules to be applied by fraud detection/classification unit 36 and stored in an ML rules database 58. As discussed in further detail below, the rules may be used to detect and/or classify a single type or category of fraudulent activity, or may be used broadly in connection with multiple types or categories of fraudulent activity. ML rule generator 40 may implement any suitable type or types of machine learning. For example, ML rule generator 40 may implement supervised learning techniques, such as decision trees, regression-based models, support vector machines (SVMs) and/or neural networks, and/or unsupervised learning techniques such as Dirichlet process mixture models and/or k-means clustering. Other machine learning techniques are also possible, such as techniques utilizing Bayesian networks, “deep learning” techniques, and so on. While shown in FIG. 1 as a single entity within AFSS 12, it is understood that ML rules database 58 may, in some embodiments, be distributed across multiple databases and/or multiple physical/hardware memories, and/or may be wholly or partially external to (e.g., remote from) AFSS 12.

External data collection unit 42 may generally collect, via network interface 32 and/or from sources internal to AFSS 12, information from various sources (e.g., FAMS 14, cardholder computing devices 20, other sources 24, etc.), and provide that data to other portions of AFSS 12 as needed (e.g., to ML rule generator 40 to generate and/or update rules, and/or to behavior analysis unit 44, dispute resolution unit 46, chargeback analysis unit 50, image analysis unit 52 and/or classification unit 54 to detect and/or classify fraudulent activity). Some data may be collected indirectly. For example, FAMS 14 may collect transaction data from merchant computing systems 22 (and/or from acquiring banks associated with one or more of merchant computing systems 22), and external data collection unit 42 may then collect that data from the account records database 30 of FAMS 14.

Once an initial set of rules has been generated and stored in ML rules database 58, those rules may dictate some or all of the types of data gathered by external data collection unit 42. In some embodiments, however, external data collection unit 42 collects a broad set of data types that may or may not be relevant to fraud determination or classification, and ML rule generator 40 continually analyzes that data to determine which data types are most predictive of fraud and/or fraud type/class.

Behavior analysis unit 44 may generally analyze cardholder-related (or other customer-related) information to identify patterns of behavior, which may then be used by

fraud detection/classification unit 36 to detect and/or classify fraudulent activity. For example, behavior analysis unit 44 may analyze information obtained from account records database 30 to identify spending patterns associated with different cardholders. The operation of behavior analysis unit 44, including the types of information analyzed and the ways in which that information is used to arrive at a result (e.g., a pattern of behavior), may be dictated by the rules stored in ML rules database 58.

Data indicative of the behavior patterns identified by behavior analysis unit 44 may be stored in an account holder behaviors database 60, for example. While shown in FIG. 1 as a single entity within AFSS 12, it is understood that account holder behaviors database 60 may, in some embodiments, be distributed across multiple databases and/or multiple physical/hardware memories, and/or may be wholly or partially external to (e.g., remote from) AFSS 12. In one embodiment, for example, account holder behaviors database 60 may be included within account records database 30. In still other embodiments, the environment 10 may not include account holder behaviors database 60, and behavior patterns may be only identified by behavior analysis unit 44 “on the fly” as needed by fraud detection/classification unit 36 (e.g., when needed to analyze a transaction in view of past spending patterns of a particular cardholder, etc.).

In some embodiments, behavior analysis unit 44 may separately analyze the transactions associated with each account holder, even if more than one account holder exists for a particular account. For example, behavior analysis unit 44 may independently analyze the transactions of each cardholder for a credit or debit card account in which each spouse has been issued a credit or debit card in his or her name. Fraud detection/classification unit 36 may then utilize the individual spending patterns when detecting and/or classifying fraud. In one embodiment where fraud detection/classification unit 36 utilizes a dollar amount threshold to detect likely fraudulent transactions, for example, a first threshold may be used for transactions made by a first cardholder listed on an account, and a higher, second threshold may be used for transactions made by a second cardholder listed on the account. Further examples are provided below in connection with FIG. 6, according to various embodiments. In this manner, fraud detection and/or classification may be made more precise than would be the case if spending patterns were only identified at the aggregate level (e.g., using a single dollar amount threshold, regardless of which cardholder made a particular transaction).

Dispute resolution unit 46 may generally analyze financial transaction data and/or other information to automatically generate queries for cardholders or other customers. For example, dispute resolution unit 46 may analyze information obtained from account records database 30. The generated queries may be designed to help fraud detection/classification unit 36 determine whether a particular transaction was fraudulent, or estimate a probability that the transaction was fraudulent, etc. Dispute resolution unit 46 may also process responses from cardholders/customers, and automatically generate additional queries based upon those responses. Examples of the operation of dispute resolution unit 46 are provided below in connection with FIGS. 4E and 9, according to various embodiments.

Chargeback analysis unit 50 may generally analyze financial transaction and/or other information to identify transactions that are good candidates for chargeback payments. For example, chargeback analysis unit 50 may analyze information obtained from account records database 30 to determine whether there is a relatively high probability that

11

the merchant (or an acquiring bank) should be responsible for a chargeback payment to a card issuer associated with FRMS 14. The operation of chargeback analysis unit 50, including the types of information analyzed and the ways in which that information is used to arrive at a result (e.g., flagging a transaction as a chargeback candidate), may be dictated by the rules stored in ML rules database 58. ML rule generator 40 may make use of chargeback rules obtained from a card network entity (e.g., from card network computing system 16), and stored in chargeback rules database 62, to generate and/or update the rules applied by chargeback analysis unit 50. Examples of the operation of chargeback analysis unit 50 are provided below in connection with FIGS. 4B and 7, according to various embodiments.

In some embodiments, transactions flagged by chargeback analysis unit 50 are subject to further, manual review using the chargeback rules stored in chargeback rules database 62. In other embodiments, chargeback analysis unit 50 (or another component of fraud detection/classification unit not shown in FIG. 1) automatically, with little or no manual input/assistance, applies the chargeback rules from chargeback rules database 62 for each flagged transaction. While shown in FIG. 1 as a single entity within AFSS 12, it is understood that chargeback rules database 62 may, in some embodiments, be distributed across multiple databases and/or multiple physical/hardware memories, and/or may be wholly or partially external to (e.g., remote from) AFSS 12.

Image analysis unit 52 may generally analyze image data corresponding to physical documents to identify fraudulent (e.g., counterfeit and/or forged) documents, and/or to flag potentially fraudulent documents for further (e.g., manual) review. For example, image analysis unit 52 may analyze information obtained from merchant computing systems 22 to determine whether there is a relatively high probability that documents presented to the merchants (e.g., personal checks, identification cards, etc.) are fraudulent. Image analysis unit 52 may be configured to analyze only a single type of document, or multiple types of documents. The operation of image analysis unit 52, including the image characteristics analyzed and the ways in which the characteristics may be used to arrive at a result (e.g., flagging a document as potentially fraudulent), may be dictated by the rules stored in ML rules database 58. Examples of the operation of image analysis unit 52 are provided below in connection with FIGS. 4F and 10, according to various embodiments.

Classification unit 54 may generally analyze broad categories of data from various sources (e.g., account records database 30, cardholder computing devices 20, merchant computing systems 22, and/or other sources 24) to categorize/classify types of suspected fraudulent financial activity. Classification unit 54 may classify fraudulent activity only within a particular subset of fraudulent financial activity (e.g., classifying debit and/or credit card transactions as involving a potential case of counterfeiting, skimming, lost/stolen card use, chargeback fraud, etc.), or may classify fraudulent financial activity across a broader spectrum (e.g., including types of identity theft not necessarily tied to a single financial transaction, such as application fraud). In some embodiments, classification unit 54 classifies suspected fraudulent activity in connection with a particular account or transaction in response to being notified of suspect activity (e.g., notified by another component of fraud detection/classification unit 36, or by a manual user input, etc.). In other embodiments, classification unit 54 identifies suspect activity before classification unit

12

54 classifies that activity. Examples of the operation of classification unit 54 are provided below in connection with FIGS. 4C and 11, according to various embodiments.

Notification unit 56 may generally provide alerts, confirmations, and/or other notifications to various individuals (e.g., customers, bank employees associated with FAMS 14, third party employees associated with AFSS 12, etc.). For example, notification unit 56 may generate a notification message stating that a fraud alert associated with a particular transaction is a false positive, and cause network interface 32 to send the message to a computer terminal or to FAMS 14 for display to a system user. As another example, notification unit 56 may cause network interface 32 to send other flagged transactions and/or documents (e.g., chargeback candidates identified by chargeback analysis unit 50, documents that image analysis unit 52 has identified as potentially fraudulent, etc.) to a computer terminal or FAMS 14 for display to a system user. As yet another example, notification unit 56 may cause network interface 32 to send queries generated by dispute resolution unit 46 to various ones of cardholder computing devices 20 for display to cardholders.

The operation of various components of the environment 10 shown in FIG. 1, according to different embodiments and/or scenarios, will be described further below in connection with the remaining figures.

III. Exemplary Process Flows for Machine Learning of Fraud Detection and/or Classification Rules

As discussed above, ML rule generator 40 may generate and/or update rules that are used for one or more of a variety of different purposes relating to fraud detection and/or classification. FIG. 2 depicts one generalized, example process flow 80 for machine learning that may be implemented by ML rule generator 40, and possibly one or more other components of fraud detection/classification unit 36.

In the process flow 80, multi-account data 82 may represent data associated with multiple financial accounts, each with one or more account holders. The financial accounts may be existing or potential accounts, and the account holders may include holders of accounts and/or potential holders of potential accounts. For example, the multi-account data 82 may include existing and/or applied-for credit card accounts, debit card accounts, savings accounts, checking accounts, investment accounts, loan accounts, etc.

Depending upon the embodiment, the multi-account data 82 may include one or more different types of information obtained (e.g., by external data collection unit 42 of FIG. 1) from one or more of FAMS 14, cardholder computing devices 20, merchant computing systems 22, and/or other sources 24. For example, the multi-account data 82 may include transaction data (e.g., transaction dates, amounts, locations, etc.) from account records database 30 of FAMS 14, data indicative of Internet Protocol (IP) addresses of cardholder computing devices 20 and/or devices in merchant computing systems 22, Internet browsing and/or search history data from cardholder computing devices 20 (or from an ISP computer system included in other sources 24, etc.), vehicle telematics data from telematics systems of cardholder vehicles, home occupancy and/or usage data (e.g., smart appliance data) from smart home systems of cardholders, autonomous or smart vehicle data, vehicle navigation system data, mobile device data, mobile device and/or vehicle GPS data, and/or one or more other types of data. In some embodiments, the multi-account data 82 only includes data that account holders or potential account holders have expressly consented to share with an entity associated with FAMS 14 and/or AFSS 12 (e.g., in exchange for fraud

protection services). In certain other embodiments, however, express consent is only needed for certain types of information, such as browsing history information, vehicle telematics data, etc.

The multi-account data **82** may be associated with multiple fraud determination labels. The labels may simply reflect whether or not fraud existed (e.g., “fraud” or “no fraud”), or may also indicate a type or class of fraud (e.g., “counterfeiting,” “lost or stolen card use,” etc.), for example. In one embodiment, each of a number of data sets in the multi-account data **82** is associated with such a label, and includes data relating to a particular financial transaction, financial account, loan application, etc., for which the fraud determination was made (e.g., after a manual and/or automated fraud investigation). The labels may include final fraud determinations that were made via earlier iterations of the process flow **80**, and/or external to the process flow **80**.

To provide a more detailed example, a first data set associated with a “card present” credit card transaction may include data describing that transaction (e.g., from account records database **30**) and data indicative of the cardholder’s online browsing activity (e.g., from one of cardholder computing devices **20**) for the 15 days immediately preceding the transaction, and be labeled “confirmed fraud.” A second data set, associated with another “card present” transaction (for the same account, or for a different account), may include the same general types of data but be labeled “no fraud,” and so on. In some embodiments and/or scenarios, the same data may appear in, or be used by, two or more of the data sets. If the two “card present” transactions described above are both associated with the same account, for example, and if the second transaction occurred less than 15 days after the first transaction, some of the same online activity data may be shared by the first and second data sets.

At a process stage **84**, the multi-account data **82** may be analyzed to generate fraud detection and/or classification rules (e.g., to be stored in ML rules database **58**). Any suitable type of supervised machine learning program/technique(s) may be used, such as SVMs, neural networks, logistic regression, etc. Generally, process stage **84** may serve to identify which type(s) of data is/are probative of whether fraud has occurred (and/or the type/category of fraud that may have occurred), and to determine the data values and/or combinations that are probative of whether fraud has occurred (and/or the type/category of fraud that may have occurred). By analyzing many (e.g., thousands) of positively and negatively labeled data sets in the multi-account data **82**, for example, process stage **84** may learn that certain spending patterns within a threshold time of a transaction tend to indicate that the cardholder made the transaction (e.g., thereby indicating that fraud has not occurred, or that a fraud report is itself fraudulent or mistaken, etc.), that certain types of online searches by a cardholder (e.g., including a descriptor of a product purchased in the transaction, or a name of the merchant, etc.) tend to indicate that the cardholder made the transaction, that the cardholder’s distance from the site of a “card present” transaction (e.g., as determined from GPS information provided by the cardholder’s smartphone, wearable electronics, or vehicle) relates to the probability of fraudulent activity according to a particular equation, and so on. Other specific examples of such rules, and how those rules may be generated, are discussed below in connection with FIGS. 3A-3F and 4A-4F, according to various embodiments.

At process stage **86**, the rules generated or updated at process stage **84** may be applied to first account data **90** associated with a particular account and customer(s) (e.g., a

customer associated with a particular one of computing devices **20**). The types of data included in first account data **90** may depend upon which types of data were determined, by process stage **84**, to be relevant to a fraud determination. For example, if the rules give weight to the amount and date of a financial transaction when determining whether the transaction is fraudulent, and also give weight to whether the account holder visits a particular type of website, then the first account data **90** may include the amount and date of one or more transactions, as well as data indicative of visited websites (e.g., Uniform Resource Locators (URLs) and/or content of visited websites, etc.). The first account data **90** may include information obtained (e.g., by external data collection unit **42**) from one or more of FAMS **14**, one of cardholder computing devices **20** associated with the customer holding the first account, one or more of merchant computing systems **22**, and/or one or more of other sources **24**, for example.

Process stage **86** may output various different types of information, depending upon the embodiment and/or scenario. For example, depending upon the content of first account data **90** and the rules generated or updated at process stage **84**, process stage **86** may generate data indicating that a particular financial transaction associated with first account data **90** is, or is not, fraudulent or potentially fraudulent. Alternatively, or additionally, process stage **86** may generate data indicating a particular classification for fraudulent or suspected fraudulent activity (e.g., a fraudulent transaction) associated with first account data **90**.

In some embodiments, further analysis (e.g., a manual review, or further automated review using additional data sources, etc.) may be performed at an additional stage, shown in dashed lines in FIG. 2 as process stage **92**. The additional analysis may then be used to make a final fraud determination (e.g., a final decision on whether fraud occurred, and/or on the type of fraud that occurred) at process stage **94**. In other embodiments, process stage **92** is omitted from process flow **80**, and process stage **94** merely represents the output of process stage **86**. The final determination made at process stage **94**, along with the first account data **90** used to make that determination, may be fed back into process stage **84** to provide additional labeled data for purposes of updating the rules.

In some embodiments, the process flow **80** includes more, fewer and/or different stages, such as any of those discussed elsewhere herein (e.g., in connection with FIGS. 3A-3F). In one alternative embodiment, process stages **84** and **86** may be combined. For example, the multi-account data **82** may be unlabeled rather than labeled (or the labels may be ignored), and the combined process stage **84**, **86** may use unsupervised learning techniques (e.g., clustering techniques) to classify anomalous/outlier financial transactions, accounts, applications, etc., as “suspect” and needing further analysis.

More specific, machine learning-based process flows generally corresponding to process flow **80** of FIG. 2 will now be described with reference to FIGS. 3A-3F. It is noted, however, that other process flows are also within the scope of the invention described herein. Moreover, while FIGS. 3A-3F generally correspond to embodiments in which supervised machine learning techniques are used, other embodiments may instead use unsupervised machine learning techniques, as noted above. In various different embodiments, fraud detection/classification unit **36** may be configured to implement only one of the process flows of FIGS. 3A-3F, or may be configured to implement two or more (e.g., all) of the process flows shown in FIGS. 3A-3F.

15

A. Exemplary Process Flow for Machine Learning of Fraud Detection Rules Using Online Activity Data

Referring first to FIG. 3A, an exemplary process flow **100** may generally be used to detect fraud using customer online activity data. In the process flow **100**, multi-customer online activity data **102** may represent data associated with the online activities of a number (e.g., thousands) of customers (e.g., credit or debit cardholders, checking or saving account holders, etc.). The multi-customer online activity data **102** may include data indicating actions that the customers took, and/or web sites visited by the customers, while the customers were connected to the Internet via web browsers (e.g., executing on respective ones of cardholder computing devices **20**). For example, the multi-customer online activity data **102** may include URLs of, and/or content (e.g., text) within, web sites visited by customers, search terms entered by customers using search engine tools, search results presented to customers by search engine tools, indications of interactive controls (e.g., virtual buttons) selected by customers on various web pages, and so on.

The multi-customer online activity data **102** may include data obtained (e.g., by external data collection unit **42** of FIG. 1) from cardholder computing devices **20**, from one or more ISPs of other sources **24**, and/or from a third party aggregator of such information, for example. In some embodiments, the multi-customer online activity data **102** may only include data that customers have expressly consented to share with an entity associated with FAMS **14** and/or AFSS **12** (e.g., in exchange for fraud protection services or other benefits, such as discounts).

As described above in connection with multi-account data **82** of process flow **80**, the multi-customer online account data **102** may be associated with multiple fraud determination labels. In some embodiments, each label may be associated with a data set that includes not only the corresponding portion of multi-customer online activity data **102**, but also one or more other types of data, such as transaction data (e.g., transaction dates, amounts, locations, etc.) for each customer from account records database **30** of FAMS **14**, data indicative of IP addresses of cardholder computing devices **20** and/or devices in merchant computing systems **22**, Internet browsing and/or search history data from cardholder computing devices **20** (or from an ISP computer system included in other sources **24**, etc.), vehicle telematics data from telematics systems of other sources **24**, home occupancy and/or usage data (e.g., smart appliance data) from smart home systems of other sources **24**, and so on. The labels may include final fraud determinations that were made via earlier iterations of the process flow **100**, and/or external to the process flow **100**. Multi-customer online account data **102** may include many (e.g., thousands) of positively and negatively labeled data sets.

At a process stage **104**, the multi-customer online activity data **102** may be analyzed to generate fraud detection rules (e.g., to be stored in ML rules database **58**). As described above in connection with process stage **84** of process flow **80**, any suitable type of supervised machine learning program/technique(s) may be used. Generally, process stage **104** may serve to identify which type(s) of online activity data is/are probative of whether fraud has occurred, and to determine the data values and/or combinations that are probative of whether fraud has occurred. While not shown in FIG. 3A, the fraud detection rules may not only detect fraud, but also classify fraud (e.g., as described below in connection with FIG. 3C), in some embodiments.

At process stage **106**, the rules generated or updated at process stage **104** may be applied to first customer online

16

activity data **110**. The first customer online activity data **110** may be associated with a particular customer, such as a customer associated with a particular one of computing devices **20**, for example. The types of data included in first customer online activity data **110** may depend upon which types of online activity data were determined, by process stage **104**, to be relevant to a fraud determination. For example, the first customer online activity data **110** may include information obtained (e.g., by external data collection unit **42**) from one of cardholder computing devices **20** (i.e., the device associated with the first customer), and/or from an ISP of other sources **24**. Some specific examples of rules that may be generated by process stage **104**, and applied at process stage **106**, are described below in connection with FIG. 4A.

Process stage **106** may output various different types of information, depending upon the embodiment and/or scenario. For example, depending upon the content of first customer online activity data **110** and the rules, process stage **106** may generate data indicating that a particular financial transaction associated with the first customer is, or is not, fraudulent or potentially fraudulent. Alternatively, or additionally, process stage **106** may generate data indicating a particular classification of fraudulent or potentially fraudulent activity associated with first customer online activity data **110**.

In some embodiments, further analysis (e.g., a manual review, or further automated review using additional data sources, etc.) is performed at an additional stage, shown in dashed lines in FIG. 3A as process stage **112**. The additional analysis may then be used to make a final fraud determination (e.g., a final decision on whether fraud occurred, and/or on the type of fraud that occurred) at process stage **114**. In other embodiments, process stage **112** is omitted from process flow **100**, and process stage **114** merely represents the output of process stage **106**.

The final determination made at process stage **114**, along with the first customer online activity data **110** (and any other data) used to make that determination, may be fed back into process stage **104** to provide additional labeled data for purposes of updating the rules. In some embodiments, a preliminary fraud determination made at process stage **106** is also fed back into process stage **104**, to allow the machine learning program to determine and improve upon past performance/accuracy.

B. Exemplary Process Flow for Machine Learning of Chargeback Candidate Detection Rules

Referring next to FIG. 3B, an exemplary process flow **120** may generally be used to identify the financial transactions for which chargebacks (e.g., post-transaction payments from merchants, or acquiring/merchant banks, back to the issuer to return proceeds from transactions) are appropriate. In the process flow **120**, multi-account transaction data **122** may represent data associated with the financial transactions involving the accounts of a number (e.g., thousands) of credit or debit cardholders. The multi-account transaction data **122** may include information such as transaction dates, transaction amounts, merchant names (and/or aliases) associated with the transaction, information relating to how the card information was collected by the merchant (e.g., by swiping, an EMV chip reader, manual entry of the card number, etc.), geographic locations of "card present" transactions, and so on. The multi-account transaction data **122** may include data obtained (e.g., by external data collection unit **42** of FIG. 1) from merchant computing systems **22** and/or from acquiring/merchant banks associated with those merchants, for example.

17

Similar to the labels described above in connection with multi-account data **82** of process flow **80**, the multi-account transaction data **122** may be associated with multiple chargeback outcome labels. For example, each label may be associated with a data set that includes the corresponding portion of multi-account transaction data **122**. The outcome labels may include final chargeback determinations that were made (in connection with the transactions represented in multi-account transaction data **122**) via earlier iterations of the process flow **120**, and/or external to the process flow **120**. Multi-account transaction data **122** may include many (e.g., thousands) of positively and negatively labeled data sets.

At a process stage **124**, the multi-account transaction data **122** may be analyzed to generate chargeback candidate detection rules (e.g., to be stored in ML rules database **58**). As described above in connection with process stage **84** of process flow **80**, any suitable type of supervised machine learning program/technique(s) may be used. Generally, process stage **124** may serve to identify which type(s) of transaction data is/are probative of whether, under the full chargeback rules of the card network entity, a chargeback is appropriate for a given transaction. Process stage **124** may also determine the transaction data values and/or combinations that are probative of whether a chargeback is appropriate for the transaction.

At a process stage **126**, the rules generated or updated at process stage **124** may be applied to first account transaction data **130** to determine whether a transaction associated with the first account is a “good” chargeback candidate. Put differently, process stage **126** may, instead of applying the full chargeback rules of the card network entity (which may be quite lengthy and complex) to the facts surrounding the transaction, use various factors and algorithms developed at process stage **124** to determine whether there exists a relatively high probability that a chargeback would be appropriate for the transaction if the full chargeback rules were applied. The process stage **126** may calculate a percentage probability that the transaction is one in which a chargeback is appropriate, for example.

The first account transaction data **130** may be associated with the account of a particular cardholder or cardholders, such as a cardholder associated with a particular one of cardholder computing devices **20**, for example. The types of data included in first account transaction data **130** may depend upon which types of transaction-related data were determined, by process stage **124**, to be relevant to a chargeback candidate determination. For example, the first account transaction data **130** may include information obtained (e.g., by external data collection unit **42**) from one of merchant computing systems **22** (e.g., the computing system of the merchant involved in the transaction being analyzed) and/or from an acquiring/merchant bank associated with that merchant. The first account transaction data **130** may also include information about one or more other transactions associated with the first account (e.g., data pertaining to other transactions occurring shortly before and/or after the transaction at issue). Some specific examples of rules that may be generated by process stage **124**, and applied at process stage **126**, are described below in connection with FIG. **4B**.

Process stage **126** may output information indicating whether the particular transaction represented by first account transaction data **130** is a “good” candidate for chargeback detection. For example, process stage **126** may output a percentage probability, calculated according to the rules generated or updated at process stage **124**, that the

18

transaction is one in which a chargeback is appropriate. As another example, process stage **126** may output a binary indicator of whether the transaction is, or is not, a strong/likely chargeback candidate (e.g., by comparing the percentage probability to a threshold probability).

If the transaction is identified as a chargeback candidate at process stage **126**, the full chargeback rules of the card network entity may be applied at a process stage **132**. Process stage **132** may include manual application of the full chargeback rules, and/or automated application of the full chargeback rules, in various different embodiments. Based upon the analysis at process stage **132**, a final chargeback determination may be made at a process stage **134**. The final determination made at process stage **134**, along with the first account transaction data **130** (and any other data) used to make that determination, may be fed back into process stage **124** to provide additional labeled data for purposes of updating the rules. In some embodiments, the indication of whether the transaction is a good chargeback candidate generated at process stage **126** may also be fed back into process stage **124**, to allow the machine learning program to determine and improve upon past performance/accuracy.

C. Exemplary Process Flow for Machine Learning of Fraud Classification Rules

Referring now to FIG. **3C**, an exemplary process flow **140** may generally be used to classify instances of suspected or potential fraud. For example, the process flow **140** may represent ongoing, real-time or batch processing of a large amount of data associated with a large number of potential and/or existing financial accounts (e.g., all accounts associated with a particular bank, or all accounts opting in to a fraud protection program, etc.). In this manner, the process flow **140** may be used to initially flag situations for closer investigation, and provide one or more classifications of the type(s) of fraud potentially at issue in order to narrow or otherwise facilitate the investigation. In other embodiments, the process flow **140** may be used to provide a narrower classification (e.g., “skimming”) when a broader class of fraud (e.g., credit card fraud) is already suspected.

In the process flow **140**, multi-account data **142** may represent data associated with financial accounts of a number (e.g., thousands) of account holders. The financial accounts may be existing or potential accounts, and the account holders may include holders of accounts and/or potential holders of potential accounts. For example, the multi-account data **142** may include existing and/or applied-for credit card accounts, debit card accounts, savings accounts, checking accounts, investment accounts, loan accounts, etc.

Depending upon the embodiment, the multi-account data **142** may include one or more different types of information obtained (e.g., by external data collection unit **42** of FIG. **1**) from one or more of FAMS **14**, cardholder computing devices **20**, merchant computing systems **22**, and/or other sources **24**. For example, the multi-account data **142** may include transaction data (e.g., transaction dates, amounts, locations, etc.) from account records database **30** of FAMS **14**, data indicative of IP addresses of cardholder computing devices **20** and/or devices in merchant computing systems **22**, Internet browsing and/or search history data from cardholder computing devices **20** (or from an ISP computer system included in other sources **24**, etc.), vehicle telematics data from telematics systems of cardholder vehicles, home occupancy and/or usage data (e.g., smart appliance data) from smart home systems of cardholders, and/or one or more other types of data. Some or all data within multi-account data **142** may be information that account holders or poten-

tial account holders have expressly consented to share with an entity associated with FAMS 14 and/or AFSS 12 (e.g., in exchange for fraud protection services).

The multi-account data 142 may be associated with multiple fraud determination labels, each indicating a type or class of fraud (e.g., “counterfeiting,” “lost or stolen card use,” “skimming,” “chargeback fraud,” “application fraud,” etc.), or indicating a lack of fraud, for example. In one embodiment, each of a number of data sets in the multi-account data 142 is associated with at least one such classification/label, and includes data relating to a particular financial transaction, financial account, loan application, etc., for which the fraud classification or classifications was/were made (e.g., after a previous iteration of process flow 140, or after another manual and/or automated fraud investigation). Multi-account data 142 may include many (e.g., thousands) of data sets labeled with various known fraud classifications.

At a process stage 144, the multi-account data 142 may be analyzed to generate fraud classification rules (e.g., to be stored in ML rules database 58). As described above in connection with process stage 84 of process flow 80, any suitable type of supervised machine learning program/technique(s) may be used. Generally, process stage 144 may serve to identify which type(s) of transaction data is/are probative of the particular type of fraud (if any) that has occurred. Process stage 144 may also determine the data values and/or combinations that are probative of the particular type of fraud (if any) that has occurred.

At a process stage 146, the rules generated or updated at process stage 144 may be applied to first account data 150. The first account data 150 may be associated with a particular account and a particular customer (e.g., a cardholder associated with a particular one of computing devices 20). The types of data included in first account data 150 may depend upon which types of data were determined, by process stage 144, to be relevant to fraud classification. For example, the first account data 150 may include information obtained (e.g., by external data collection unit 42) from one or more of FAMS 14, one of cardholder computing devices 20 (i.e., the device associated with the customer holding or applying for the first account), one or more of merchant computing systems 22, and/or one or more of other sources 24. Some specific examples of rules that may be generated by process stage 144, and applied at process stage 146, are described below in connection with FIG. 4C.

Process stage 146 may output data (e.g., a message or code) that is used to classify suspected fraudulent activity (in connection with the account associated with first account data 150) at a process stage 152. For example, process stage 152 may assign a classification of “counterfeiting” if process stage 146 determined that the first account data 150 indicated a number of circumstances that, according to the rules generated at process stage 144, are known to be correlated with counterfeiting activity (e.g., two “card present” transactions occurring in different states within the same one-hour time period, etc.). In some embodiments and/or scenarios, two or more classifications may concurrently be assigned to first account data 150. For example, process stage 146 may determine a set of probabilities for a set of two or more potential types of fraud, and process stage 152 may assign each classification, with each respective probability, to first account data 150. Moreover, in some embodiments and scenarios, process stage 152 may assign a classification that corresponds to an absence of any suspected fraud (e.g., “no fraud”).

At a process stage 154, if process stage 152 assigned a classification other than one indicating the absence of suspected fraud, the first account data 150, and/or other information associated with the account and the suspected class of fraud, may be analyzed in depth to make a final fraud determination at a process stage 156. Generally, the fraud classification may be used to facilitate the analysis at process stage 154, with process stage 154 including manual and/or automated fraud detection techniques. For example, personnel associated with AFSS 12 may use the fraud classification(s) to inform their strategy and/or focus with respect to conducting an in-depth fraud investigation.

The additional analysis at process stage 154 may then result in a final fraud determination at process stage 156. The final determination may indicate both whether fraud occurred and, if so, the class(es)/type(s) of fraud that occurred. The final determination made at process stage 156, and information used to make that determination (e.g., the first account data 150 and potentially other data), may be fed back into process stage 144 to provide additional labeled data for purposes of updating the rules. In some embodiments, the (preliminary) fraud classification made at process stage 152 may also be fed back into process stage 144 to help the machine learning program identify instances in which the preliminary classifications at process stage 152 were incorrect. Process stage 144 may then update the fraud classification rules in ways that seek to prevent or reduce such instances in the future.

D. Exemplary Process Flow for Machine Learning of Application Fraud Detection Rules

Referring now to FIG. 3D, an exemplary process flow 160 may generally be used to detect application fraud. “Application fraud” may generally refer to fraud in connection with the application for any type of financial account, loan and/or line of credit (e.g., mortgage loan, vehicle loan, small business loan, payday loan, home equity line of credit, credit card account, debit card account, checking account, savings account, investment account, etc.). In some embodiments and/or scenarios, however, the application may be for non-financial purposes, such as an application for membership in a particular group or institution, for example.

In the process flow 160, multi-applicant search history data 162 may represent data associated with the Internet search history of a number (e.g., thousands) of applicants. The multi-applicant search history data 162 may include search terms entered by the applicants using online search engine tools, for example, and/or the results of such searches (e.g., URLs, titles and/or contents of search results), for example.

The multi-applicant search history data 162 may include data obtained (e.g., by external data collection unit 42 of FIG. 1) from cardholder computing devices 20, from one or more ISPs of other sources 24, and/or from a third party aggregator of such information, for example. In some embodiments, the multi-applicant search history data 162 only includes data that the applicants have expressly consented to share with an entity associated with FAMS 14 and/or AFSS 12 (e.g., in exchange for consideration of their applications).

As described above in connection with multi-account data 82 of process flow 80, the multi-applicant search history data 162 may be associated with multiple fraud determination labels. In some embodiments, each label may be associated with a data set that corresponds to an application submitted by a particular applicant, where the data set includes the corresponding portion of multi-applicant search history data 162 (e.g., the search terms and/or results associated with the

21

particular application). The labels may include final fraud determinations that were made via earlier iterations of the process flow **160**, and/or external to the process flow **160**. Multi-applicant search history data **162** may include many (e.g., thousands) of positively and negatively labeled data sets.

At a process stage **164**, the multi-applicant search history data **162** may be analyzed to generate application fraud detection rules (e.g., to be stored in ML rules database **58**). As described above in connection with process stage **84** of process flow **80**, any suitable type of supervised machine learning program/technique(s) may be used. Generally, process stage **164** may serve to identify which type(s) of Internet search-related data is/are probative of whether application fraud has occurred, and to determine the data values and/or combinations that are probative of whether application fraud has occurred.

At process stage **166**, the rules generated or updated at process stage **164** may be applied to first applicant search history data **170**. The first applicant search history data **170** may be associated with a particular application and a particular applicant (e.g., a person associated with a particular one of computing devices **20**), for example. The types of data included in first applicant search history data **170** may depend upon which types of Internet search-related data were determined, by process stage **164**, to be relevant to a fraud determination. The first applicant search history data **170** may include information obtained (e.g., by external data collection unit **42**) from one of computing devices **20** (i.e., the device associated with the first applicant), and/or from an ISP of other sources **24**, for example. Some specific examples of rules that may be generated by process stage **164**, and applied at process stage **166**, are described below in connection with FIG. **4D**.

Process stage **166** may output information indicating whether fraud is suspected in connection with the application corresponding to first applicant search history data **170**. For example, process stage **166** may output a percentage probability, calculated according to the rules generated or updated at process stage **164**, that the application was fraudulently made (e.g., by someone other than the purported applicant or an authorized representative thereof). As another example, process stage **166** may output a binary indicator of whether the application likely was, or likely was not, fraudulently made (e.g., by comparing a percentage probability to a threshold probability).

In some embodiments, further analysis (e.g., a manual review, or further automated review using additional data sources, etc.) is performed at an additional stage, shown in dashed lines in FIG. **3D** as process stage **172**. The additional analysis may then be used to make a final fraud determination (e.g., a final decision on whether application fraud occurred) at process stage **174**. In other embodiments, process stage **172** is omitted from process flow **160**, and process stage **174** merely represents the output of process stage **166**. The final determination made at process stage **174**, along with the first applicant search history data **170** (and any other data) used to make that determination, may be fed back into process stage **164** to provide additional labeled data for purposes of updating the rules. In some embodiments, a preliminary fraud determination made at process stage **166** is also fed back into process stage **164**, to allow the machine learning program to determine and improve upon past performance/accuracy.

E. Exemplary Process Flow for Machine Learning of Fraud Dispute Resolution Rules

22

Referring now to FIG. **3E**, an exemplary process flow **180** may generally be used to facilitate the resolution of fraud disputes (or potential disputes) with customers/account holders. For example, the process flow **180** may be used to determine whether a reportedly unauthorized or fraudulent transaction (e.g., one that the account holder reported as such when looking at his or her account statement) was indeed unauthorized or fraudulent. In some embodiments, the process flow **180** may also, or instead, be used to determine whether an “unrecognized” transaction (i.e., one that the account holder does not recall, but does not necessarily report as fraudulent) was unauthorized or fraudulent.

In the process flow **180**, multi-account data **182** may represent data associated with financial accounts of a number (e.g., thousands) of account holders. For example, the multi-account data **182** may include data associated with financial transactions relating to credit card accounts, debit card accounts, savings accounts, checking accounts, etc. For ease of explanation, FIG. **3E** will be described with reference to an embodiment in which the accounts are credit card accounts.

In one embodiment, the multi-account data **182** may include transaction data (e.g., transaction dates, amounts, locations, etc.) obtained from FRMS **14** (e.g., by external data collection unit **42** of FIG. **1**). In some embodiments, however, the multi-account data **182** also includes information obtained from cardholder computing devices **20**, merchant computing systems **22**, and/or other sources **24**. For example, the multi-account data **182** may include, in addition to transaction data from account records database **30** of FAMS **14**, data indicative of IP addresses of cardholder computing devices **20** and/or devices in merchant computing systems **22**, Internet browsing and/or search history data from cardholder computing devices **20** (or from an ISP computer system included in other sources **24**, etc.), vehicle telematics data from telematics systems of cardholder vehicles, home occupancy and/or usage data (e.g., smart appliance data) from smart home systems of cardholders, autonomous vehicle data, smart vehicle data, mobile device data, vehicle or mobile device GPS data, and/or one or more other types of data. Some or all data within multi-account data **182** may be information that account holders or potential account holders have expressly consented to share with an entity associated with FAMS **14** and/or AFSS **12** (e.g., in exchange for fraud protection services).

As described above in connection with multi-account data **182** of process flow **80**, the multi-account data **182** may be associated with multiple fraud determination labels (e.g., “fraud” and “no fraud,” and/or more complex labels that indicate type/class, such as “lost/stolen card use,” etc.). In some embodiments, each label may be associated with a data set that includes the corresponding portion of multi-account data **182**. The labels may include final fraud determinations that were made via earlier iterations of the process flow **180**, and/or external to the process flow **180**. Multi-account data **182** may include many (e.g., thousands) of positively and negatively labeled data sets.

At a process stage **184**, the multi-account data **182** may be analyzed to generate query generation rules (e.g., to be stored in ML rules database **58**). As described above in connection with process stage **84** of process flow **80**, any suitable type of supervised machine learning program/technique(s) may be used. Generally, process stage **184** may serve to identify which types of information are probative of whether fraud has occurred, and to craft rules that formulate queries to ascertain such information based upon account data.

23

For example, process stage **184** may determine that, for a suspect “card present” transaction, a verified, non-fraudulent “card present” transaction within 10 miles and 3 hours of the suspect transaction is probative of whether the suspect transaction was fraudulent. Based upon this finding, process stage **184** may also generate a rule specifying that a cardholder should be queried as to whether he/she can confirm making each “card present” transaction within 10 miles and 3 hours of the suspect transaction. As another example, process stage **184** may determine that a merchant using a billing alias different from its legal and/or commonly-known name (e.g., by at least some threshold level of similarity, as measured by number of similar characters, order of characters, etc.) is probative of whether the cardholder authorized a transaction associated with that billing alias. Based upon this finding, process stage **184** may generate a rule specifying that a cardholder should be queried as to whether he/she is aware of a billing alias used for a suspect transaction if that billing alias is sufficiently different from the legal/common name of the merchant.

At process stage **186**, the rules generated or updated at process stage **184** may be applied to first account data **190**. The first account data **190** may be associated with a particular cardholder, such as a cardholder associated with a particular one of cardholder computing devices **20**, for example. The types of data included in first account data **190** may depend upon which types of data were determined, by process stage **184**, to be relevant to developing dispute resolution queries. Process stage **186** may generate a set of one or more queries in accordance with the rules and the contents of first account data. Some specific examples of rules that may be generated by process stage **184** and applied at process stage **186**, and the queries that may be generated as a result, are described below in connection with FIG. 4E.

At a process stage **192**, the generated queries may be sent to the cardholder in one or more of various ways, such as sending the queries via SMS text message and/or email, and/or via a web browser or dedicated application executing on the one of cardholder computing devices **20** that is associated with the cardholder, for example. At a process stage **194**, responses to the queries are received from the cardholder (e.g., via inputs made by the cardholder via the web browser or application, or a responsive SMS text message or email, etc.). In some embodiments, the rules generated or updated at process stage **184** specify the manner in which follow-up queries should be generated based upon the responses received at process stage **194**, and process stages **192** and **194** may be repeated multiple times.

In some embodiments, further analysis (e.g., a manual review, or further automated review using additional data sources, etc.) that makes use of the received responses is performed at an additional stage, shown in dashed lines in FIG. 3E as process stage **196**. The additional analysis may then be used to make a final fraud determination (e.g., a final decision on whether fraud occurred, and/or on the type of fraud that occurred) at process stage **198**. In other embodiments, process stage **196** is omitted from process flow **180**, and process stage **198** is based upon information from the cardholder. For example, the questions generated at process stage **192** may “jog” the cardholder’s memory, and cause him or her to indicate that the transaction at issue was authorized. The final determination made at process stage **198**, along with the first account data **110** (and any other data used at process stage **196**), the queries generated at process stage **186** and/or the responses received at process stage **194**, may be fed back into process stage **184** to provide additional labeled data for purposes of updating the rules.

24

F. Exemplary Process Flow for Machine Learning of Document Fraud Detection Rules

Referring now to FIG. 3F, an exemplary process flow **200** may generally be used to detect fraud relating to documents, such as counterfeit and/or forged documents. The process flow **200** may be used in connection with various kinds of documents, such as checks (e.g., personal checks, cashier’s checks, etc.), money orders, treasury bills, identification documents (e.g., social security cards, driver’s licenses, passports, birth certificates, etc.), certification documents, and so on.

In the process flow **200**, multi-document image data **202** may represent digital images of a number (e.g., thousands) of physical documents of one or more types. The multi-document image data **202** may include images in one or more formats, such as raster formats (e.g., JPEG, TIFF, GIF, BMP, PNG, etc.) and/or vector formats (e.g., CGM, SVG, etc.), for example. The multi-document image data **202** may include data obtained (e.g., by external data collection unit **42** of FIG. 1) from merchant computing systems **22** (e.g., point-of-sale devices with cameras for document identification) and/or from FAMS **14** (e.g., images of personal checks), for example. In some embodiments, the multi-document image data **202** may only include data representing images that customers (or other individuals associated with the documents) have expressly consented to share (e.g., as a prerequisite to making a purchase, or in exchange for fraud protection services, etc.).

As described above in connection with multi-account data **82** of process flow **80**, the multi-document image data **202** may be associated with multiple fraud determination labels. In some embodiments, each label may be associated with data representing a digital image of a particular document. The labels may include final fraud determinations (e.g., “fraud” or “no fraud,” or more complex labels such as “forgery,” “counterfeit,” “forgery—signature,” “counterfeit—angular line offset(s) outside tolerance,” etc.) that were made via earlier iterations of the process flow **200**, and/or external to the process flow **200**. Multi-document image data **202** may include many (e.g., thousands) of positively and negatively labeled data sets.

At a process stage **204**, the multi-document image data **202** may be analyzed to generate document fraud detection rules (e.g., to be stored in ML rules database **58**). As described above in connection with process stage **84** of process flow **80**, any suitable type of supervised machine learning program/technique(s) may be used. Generally, process stage **204** may serve to identify which characteristics of a document are probative of whether the document is counterfeit, and to determine the ranges, tolerances, etc., that are probative of whether the document is counterfeit. In some embodiments, process stage **204** also, or instead, identifies which characteristics of information entered in document fields are probative of whether the document was forged (e.g., drafted or populated by someone other than the person purported to have drafted or populated the document).

At process stage **206**, the rules generated or updated at process stage **204** may be applied to first document image data **210**. The first document image data **210** may be digital image data corresponding to a particular, physical document. The first document image data **210** may include information obtained (e.g., by external data collection unit **42**) from one of merchant computing systems **22** (e.g., for real-time verification of an identification or other document presented during or prior to a sale), or from FAMS **14** (e.g., for real-time or batch-processing verification of a personal

25

check prior to clearing the check), for example. Some specific examples of rules that may be generated by process stage **204**, and applied at process stage **206**, are described below in connection with FIG. 4F.

Process stage **206** may output information indicating whether fraud is suspected in connection with the document corresponding to first document image data **210**. For example, process stage **206** may output two percentage probabilities calculated according to the rules generated or updated at process stage **204**, with the first indicating the likelihood that the document is counterfeit and the second indicating the likelihood that the document includes forged content. As another example, process stage **206** may output binary indicators of whether the document likely is, or likely is not, counterfeit and/or includes forged content (e.g., by comparing percentage probabilities to threshold probabilities).

In some embodiments, further analysis (e.g., a manual review, or further automated review using additional data sources, etc.) may be performed at a process stage **212**. The additional analysis may then be used to make a final fraud determination (e.g., a final decision on whether the document is fraudulent) at process stage **214**. For example, the process stage **206** may act as a filter, and flag only those documents having a relatively high probability of being fraudulent. In this manner, a considerably smaller amount of human and/or processing resources may be consumed at process stage **212**.

The final determination made at process stage **214**, along with the first document image data **210** used to make that determination, may be fed back into process stage **204** to provide additional labeled data for purposes of updating the rules. In some embodiments, a preliminary fraud determination made at process stage **206** may also be fed back into process stage **204**, to allow the machine learning program to determine and improve upon past performance/accuracy.

IV. Exemplary Rules for Fraud Detection and/or Classification

FIGS. 4A-4F depict exemplary factors and algorithms that may be used in connection with various fraud detection and/or classification rules, according to different embodiments. It is noted that the rule sets corresponding to FIGS. 4A-4F are purely for purposes of illustration and are not limiting. Particularly in embodiments where machine learning is utilized, for example, the algorithms and/or factors may be far more complex, and/or less intuitive, than some or all of the examples shown in FIGS. 4A-4F.

A. Exemplary Fraud Detection Rule Set Using Online Activity

Referring first to FIG. 4A, an exemplary rule set **220** (e.g., generated at process stage **104** of FIG. 3A) may use various factors relating to online activity of a cardholder to detect fraud in connection with a particular credit or debit card transaction. The rule set **220** may correspond to a particular embodiment and scenario in which the transaction at issue is a “card present” transaction, and in which the rule set **220** seeks to determine whether the cardholder made or otherwise authorized the transaction. The rule set **220** may be incorporated into a review process that is generally applied to all transactions, a review process applied only to those transactions that were flagged by a preliminary fraud alert, or a review process applied only after a cardholder reports the transaction as unauthorized, for example.

The factors considered under the rule set **220** may include a number of interest-based factors **222** and a number of location-based factors **224**. The interest-based factors **222** may relate to the cardholder’s interest (or non-interest) in a

26

product or service purchased via the transaction, and/or the merchant providing the product or service, while the location-based factors **224** may relate to the cardholder’s location or probable location.

As seen in FIG. 4A, the interest-based factors **222** may include: (1) whether the cardholder searched online for the specific product or service purchased via the transaction at issue (e.g., by determining whether search terms entered by the cardholder included the name of the product or service involved in the transaction, or included a description of the product or service, etc.); (2) whether the cardholder visited a website associated with the merchant (e.g., by comparing URLs of websites visited by the cardholder to a known URL of the merchant’s website, or by searching the contents of websites visited by the cardholder for the merchant’s name, etc.); (3) whether the cardholder endorsed the merchant, or the product or service provided by the merchant, via a social media account of the cardholder (e.g., by determining whether the cardholder “liked” the merchant, product or service via his or her Facebook® account, etc.); (4) whether the cardholder visited a website associated with a competitor of the merchant (e.g., by comparing URLs of websites visited by the cardholder to known URLs of known competitors’ websites, or by searching the contents of websites visited by the cardholder for the competitors’ names, etc.); (5) whether the cardholder searched online for a different product or service in the same price range as the transaction amount (e.g., by analyzing search terms and/or results, and/or by analyzing URLs or contents of websites visited by the cardholder and comparing prices of products/services, etc.); and/or (6) whether the cardholder entered search terms indicative of the cardholder’s need for the product or service (e.g., by determining that the cardholder entered search terms including “pipe leak” prior to the purchase of new plumbing hardware, or “computer repair” prior to the purchase of a new hard drive, etc.). In other embodiments, the interest-based factors **222** may include more, fewer and/or different factors than those shown in FIG. 4A.

As is also seen in FIG. 4A, the location-based factors **224** may include: (1) whether the cardholder “checked in” to a flight having a destination near the location where the transaction was initiated (e.g., by determining whether the cardholder checked in to a flight having a destination at the city in which the transaction occurred, or within a threshold number of miles of the city in which the transaction occurred, etc.); (2) whether the cardholder visited a website associated with a place near (or in) which the transaction was initiated (e.g., by comparing URLs of websites visited by the cardholder to URLs of websites known to be associated with particular areas, and/or by searching the contents of websites visited by the cardholder for location or area names, etc.); and/or (3) whether the cardholder endorsed a place near (or in) which the transaction was initiated via a social media account of the cardholder (e.g., by determining whether the cardholder “liked” the geographic area, attraction or other place via his or her Facebook® account, etc.). In other embodiments, the location-based factors **224** may include more, fewer and/or different factors than those shown in FIG. 4A.

Generally, the data indicative of whether the circumstance corresponding to each of interest-based factors **222** and/or location-based factors **224** is present/true for a particular cardholder may be included in the first customer online activity data **110** described above in connection with FIG. 3A. For example, external data collection unit **42** of FIG. 1 may obtain the search terms, URLs, user online selections, etc., needed to determine whether the various factors exist,

from the cardholder's computing device (e.g., one of cardholder computing devices **20**) and/or from an ISP of other sources **24**.

As is also seen in FIG. **4A**, each of the interest-based factors **222** and location-based factors **224** may be associated with a particular score or weighting value. In the rule set **220** shown in FIG. **4A**, a total score may be calculated based upon which factors are, or are not, present (e.g., add 94 points if it is determined that the cardholder searched for the particular lawnmower model that was purchased, add another 80 points if the transaction was a "card present" transaction in the Chicago suburb of Joliet and the cardholder checked in to a flight to Chicago just prior to the transaction, etc.).

In some embodiments, certain factors may instead be associated with negative scores (e.g., minus 80 if the cardholder checked in to a flight with a destination at least 200 miles from the site of the transaction and within one day of the transaction, etc.). Moreover, certain factors may be associated with metrics or algorithms that determine how heavily those factors are weighed. As indicated in FIG. **4A**, for example, search terms entered by the cardholder may be used to calculate a "need score" X (e.g., where X is based upon frequency of certain search terms being used, the amount of time spent clicking through search results, the magnitude and/or urgency of a problem indicated by the search terms, etc.), with X then being used to calculate a score equal to $0.2 \times$.

The rule set **220** may then output the total score (e.g., $94+80=+174$), a normalized total score, an indication of whether the total score exceeded a threshold (e.g., a threshold of $+100$), a probability calculated based upon the total score, and/or some other indicator or measure of the existence or likelihood of fraud. In the example shown in FIG. **4A**, it can be seen that larger scores generally correspond to a greater probability that the transaction was made or authorized by the cardholder. If the transaction is being automatically reviewed (e.g., to determine whether a fraud alert is appropriate, without any initial input from the cardholder), this may mean that a lower score corresponds to a higher probability of fraud. Conversely, if the cardholder had reported the transaction as being fraudulent, a higher score may correspond to a higher probability of fraud (i.e., fraud on the part of the cardholder).

In some embodiments, the rule set **220** may also include one or more other types of factors not necessarily based upon online activities of the cardholder (e.g., whether GPS of the cardholder's smartphone or vehicle indicates that he or she was in that area shortly before or after the transaction, etc.), and/or may omit either interest-based factors **222** or location-based factors **224**.

B. Exemplary Chargeback Candidate Detection Rule Set

Referring next to FIG. **4B**, an exemplary rule set **230** (e.g., generated at process stage **124** of FIG. **3B**) may use various factors relating to a transaction between a cardholder and a merchant to determine whether the transaction should be flagged as a candidate for a chargeback (e.g., to determine whether the transaction should be reviewed under a full set of chargeback rules associated with the appropriate card network entity). The rule set **230** may correspond to a particular embodiment and scenario in which the transaction at issue is a "card present" transaction.

As seen in FIG. **4B**, the factors considered under the rule set **230** may include: (1) whether an EMV chip card was not inserted in a point-of-sale EMV chip reader device of the merchant; (2) whether a non-EMV card was not swiped in a point-of-sale device of the merchant; (3) whether the card

is past its expiration date; (4) whether the transaction is for the same amount and/or date as another transaction involving the same card and merchant (e.g., by analyzing other transactions involving the same account and merchant within a particular time span); and/or (2) whether the transaction is for greater than a threshold amount. For example, one of merchant computing systems **22** of FIG. **1** (or an acquiring/merchant bank) may provide transaction details that include the amounts, dates, etc., to FAMS **14** for storage in account records database **30**, and external data collection unit **42** may then retrieve that information from account records database **30**. Generally, the data indicative of whether the circumstance corresponding to each of the factors is present/true for a particular transaction may be included in the first account transaction data **130** described above in connection with FIG. **3B**. In other embodiments, the factors considered under rule set **230** may include more, fewer and/or different factors than those shown in FIG. **4B**. It is noted that, in some embodiments, one or more factors may simply relate to the desirability (e.g., from a card issuer perspective) of further reviewing whether a chargeback is appropriate, without necessarily relating to the likelihood that a chargeback is appropriate.

As is also seen in FIG. **4B**, each of the factors may be associated with a particular score or weighting value. A total score may be calculated based upon which factors are, or are not, present (e.g., add 62 points if it is determined that the transaction has the same amount and date as another transaction occurring close in time and involving the same card and merchant). In some embodiments, certain factors may instead be associated with negative scores, and/or certain factors may be associated with metrics or algorithms that determine how heavily those factors are weighed.

The rule set **230** may then output the total score, a normalized total score, an indication of whether the total score exceeded a threshold, a probability calculated based upon the total score, and/or some other indicator or measure of the likelihood that a chargeback is appropriate for the transaction. In the example shown in FIG. **4B**, it can be seen that larger scores generally correspond to a greater probability that a chargeback is appropriate.

C. Exemplary Fraud Classification Rule Set

Referring now to FIG. **4C**, an exemplary rule set **240** (e.g., generated at process stage **144** of FIG. **3C**) may use a diverse array of factors to classify the type(s) of fraudulent activity, if any, that is/are suspected to be associated with an event or series of events. The rule set **240** may correspond to a particular embodiment and scenario in which the event at issue is a financial transaction involving a debit or credit card. In other embodiments and/or scenarios, however, the rule set **240** may classify fraudulent activity with respect to specific other types of events (e.g., loan applications), or may detect a variety of different event types (e.g., various types of financial transactions, loan or credit applications, etc.) and broadly classify fraudulent activity in connection with the detected event types (e.g., lost/stolen card use, application fraud, etc.).

In one embodiment, each potential classification (with the possible exception of "no fraud") may be associated with a number of factors probative of whether that type/class of fraud has occurred. As seen in FIG. **4C**, for example, the rule set **240** may include counterfeit factors **242** (e.g., factors indicating that a counterfeit card was used for the transaction), account takeover factors **244** (e.g., factors indicating that the transaction resulted from an unauthorized person gaining online access to the credit or debit card account itself, via phishing, malware or other means), chargeback

fraud factors **246** (e.g., factors indicating that the cardholder made or otherwise authorized a purchase that the cardholder later contested) and skimming factors **248** (e.g., factors indicating that the card information used for the transaction was obtained via a skimming card reader device illegally installed in an ATM, gas station pump or other location). In other embodiments, the rule set **240** may also, or instead, include factors corresponding to one or more other fraud classifications (e.g., forgery, lost/stolen card use, etc.).

As seen in FIG. 4C, the counterfeit factors **242** may include: (1) whether the suspect transaction and another, contemporaneous transaction (e.g., occurring within one hour, etc.) in another state are both “card present” transactions; and/or (2) if the suspect transaction is a “card present” transaction, whether the card (if an EMV chip card) was not inserted in an EMV chip card reader. For example, one or more of merchant computing systems **22** of FIG. 1 (or one or more acquiring/merchant banks) may provide transaction details that include whether the transaction was “card present,” whether the card was inserted in an EMV chip card reader, etc., to FAMS **14** for storage in account records database **30**, and external data collection unit **42** may then retrieve that information from account records database **30**. In other embodiments, the counterfeit factors **242** may include more, fewer and/or different factors than those shown in FIG. 4C.

The account takeover factors **244** may include: (1) whether the debit or credit card account password was changed within the 10 days prior to the transaction; and/or (2) whether the transaction was originated from an IP address not associated with the cardholder. For example, external data collection unit **42** may retrieve password change information from account records database **30** of FIG. 1, which may log all password update activity, and/or may retrieve IP address information from one of merchant computing systems **22** (e.g., the computing system of the merchant involved in the transaction). In other embodiments, the account takeover factors **244** may include more, fewer and/or different factors than those shown in FIG. 4C.

The chargeback fraud factors **246** may include: (1) whether the cardholder had searched online for the product or service purchased via the transaction; and/or (2) whether the cardholder had visited a website associated with the merchant involved in the transaction. For example, external data collection unit **42** of FIG. 1 may retrieve online search information (e.g., search terms and/or results) and/or URLs from the one of cardholder computing devices **20** that is associated with the cardholder, and/or from an ISP (of other sources **24**) used by the cardholder. In other embodiments, the chargeback fraud factors **246** may include more, fewer and/or different factors than those shown in FIG. 4C.

The skimming factors **248** may include: (1) the number (X) of earlier transactions in which the card used for the transaction at issue was used at an ATM machine or a gas station pump within the 10 days prior to the transaction at issue; and/or (2) whether the transaction at issue originated from an IP address not associated with the cardholder. For example, external data collection unit **42** of FIG. 1 may retrieve transaction data indicating that certain past purchases were made using gas station pump card readers, and/or indicating that the card was used for one or more ATM withdrawals, from account records database **30**, and/or may retrieve the originating IP address from the one of merchant computing systems **22** associated with the merchant involved in the transaction at issue. In other embodiments, the skimming factors **248** may include more, fewer and/or different factors than those shown in FIG. 4C.

Generally, the data indicative of whether the circumstance corresponding to each of counterfeit factors **242**, account takeover factors **244**, chargeback fraud factors **246** and/or skimming factors **248** is present/true for a particular transaction may be included in the first account data **150** described above in connection with FIG. 3C, for example.

As is also seen in FIG. 4C, each of the counterfeit factors **242**, account takeover factors **244**, chargeback fraud factors **246** and skimming factors **248** may be associated with a particular score or weighting value. The factors for each classification (counterfeit, account takeover, chargeback fraud, skimming) may be used to calculate a total score specific to that classification. In the rule set **240** shown in FIG. 4C, for example, a counterfeit score may be calculated based upon which of factors **242** are, or are not, present, an account takeover score may be calculated based upon which of factors **244** are, or are not, present, and so on. In some embodiments, certain factors may instead be associated with negative scores, and/or certain factors (e.g., the first of skimming factors **248** shown in FIG. 4C) may be associated with metrics or algorithms that determine how heavily those factors are weighed.

For each classification/category, the rule set **240** may output the total score, a normalized total score, an indication of whether the total score exceeded a threshold, a probability calculated based upon the total score, and/or some other indicator or measure of the likelihood that fraud of that particular type/class occurred in connection with the transaction. In the example shown in FIG. 4C, it can be seen that larger scores generally correspond to a greater probability that the respective classification is accurate. Referring back to FIG. 3C, the classification at process stage **152** may be the classification having the highest score and/or probability under rule set **240**, or may include the score and/or probability for each classification, the top three classifications, etc.

D. Exemplary Application Fraud Detection Rule Set

Referring now to FIG. 4D, an exemplary rule set **260** may use online search information (e.g., search terms, search results, clicked/selected search results, etc.) to detect whether an application was fraudulent (e.g., not populated and/or submitted by the purported applicant). The rule set **260** may have been generated at process stage **164** of FIG. 3D, for example. The rule set **260** may be incorporated into a review process that is generally applied to all applications received by a particular entity or anti-fraud service, or a review process applied only to those applications that were flagged by a preliminary fraud alert, for example.

The factors considered under the rule set **260** may generally be probative of whether the person that submitted the application (e.g., via a web browser, a dedicated application, as an email attachment, by snail mail, etc.) had performed one or more online searches indicating that he or she was trying to learn more about the purported applicant in order to populate particular fields of the application (e.g., a “home address” field, “employment history” fields, etc.). The “purported applicant” may be a person whose name appears in a name and/or signature field of the application, for example.

As seen in FIG. 4D, the factors of exemplary rule set **260** may include: (1) whether the applicant used search terms that included the name of the purported applicant; (2) whether the search terms also included the words “address” or “residence” (and possibly other synonyms or near-synonyms); and/or (3) whether the search terms also included the words “employer,” “job” and/or “career” (and possibly other synonyms or near-synonyms). In other embodiments, the rule set **260** may include more, fewer and/or different

factors than those shown in FIG. 4D. For example, the rule set **260** may include one or more factors relating to which search results appeared and/or were selected (e.g., “clicked” on after appearing on a user interface) by the applicant.

Generally, the data indicative of whether the circumstances corresponding to the factors of rule set **260** are present/true for a particular applicant may be included in the first applicant search history data **170** described above in connection with FIG. 3D. For example, external data collection unit **42** of FIG. 1 may obtain the search terms, search results, search result user selections, etc., needed to determine whether the various factors exist, from the applicant’s computing device (e.g., similar to one of cardholder computing devices **20**) and/or from an ISP of other sources **24**. Access to such information may be made a condition of having the application be considered, for example.

As is also seen in FIG. 4D, each of the factors of rule set **260** may be associated with a particular score or weighting value. A total score may then be calculated based upon which factors are, or are not, present. In some embodiments, certain factors may instead be associated with negative scores, and/or certain factors may be associated with metrics or algorithms that determine how heavily those factors are weighed.

The rule set **260** may then output the total score, a normalized total score, an indication of whether the total score exceeded a threshold, a probability calculated based upon the total score, and/or some other indicator or measure of the existence or likelihood of application fraud. In the example shown in FIG. 4D, it can be seen that larger scores may generally correspond to a greater probability that the application was not populated and/or submitted by the purported applicant.

E. Exemplary Fraud Dispute Resolution Rule Set

Referring now to FIG. 4E, a flow diagram illustrates at least a portion of a process flow **270** implementing an exemplary rule set for fraud dispute, or potential fraud dispute, resolution (e.g., a rule set generated at process stage **184** of FIG. 3E). The process flow **270** may be used to help resolve a dispute over a contested transaction, or to help a customer recall an unrecognized transaction, for example. FIG. 4E illustrates a process flow, rather than just a set of factors, in order to better illustrate an example process for generating queries based upon the generated rules, according to one embodiment. The process flow **270** may correspond to a particular embodiment and scenario in which the transaction subject to dispute or potential dispute is a credit or debit card transaction.

In the exemplary process flow **270**, the rule set may specify that a process stage **272** determines whether the transaction was a “card present” transaction. If not, the rule set may specify that the flow proceed directly to a process stage **280**. If so, however, the rule set may specify that the flow instead proceeds to a process stage **274**.

The rule set may also specify that process stage **274** determines whether at least one other transaction associated with the cardholder’s account occurred within some threshold number of hours (X) of the transaction at issue. If not, the rule set may specify that the flow proceeds directly to process stage **280**. If so, however, the rule set may specify that the flow instead proceeds to a process stage **276**.

Process stage **276** may generate one or more location-related queries using transaction data associated with the cardholder’s account. The queries may ask, for example, whether the cardholder was in (or near) one or more particular geographic areas or locations at various times. If the transaction at issue occurred in San Francisco, for example,

with a first other “card present” transaction occurring in Santa Rosa four hours earlier and a second other “card present” transaction occurring in San Jose two hours later, process stage **276** may generate one or more queries asking whether the cardholder made or authorized the earlier and/or later transactions, and/or whether the cardholder traveled on a route from Santa Rosa to San Jose that passed through San Francisco, etc.

In some embodiments, the location-related queries are generated based upon data associated with events or circumstances other than transactions. For example, if the transaction at issue occurred in Sarasota, Fla., and the data considered under the rule set indicates that the cardholder checked in to a flight to Tampa, process stage **276** may generate one or more queries asking whether the cardholder completed the flight, where the cardholder went after landing in Tampa, etc.

The rule set may also specify that process stage **280** determines whether the transaction at issue is associated with a billing alias that is dissimilar to the name of the merchant involved in the transaction. For example, the computing system of the merchant (e.g., one of merchant computing systems **22** of FIG. 1) may have sent to FAMS **14** a transaction record that identified the merchant by the alias, and was presented to the cardholder as an online or paper account statement. The determination at process stage **280** may use the billing alias to identify a legal and/or common name of the merchant (e.g., using a relational database stored in AFSS**12** or FAMS **14**), and determine that there is at least some threshold level of dissimilarity (e.g., based upon difference of characters, character ordering, etc.) between the billing alias and the merchant name.

If the billing alias and merchant name are not sufficiently dissimilar, the rule set may specify that the flow proceeds directly to a process stage **284**. If sufficiently dissimilar, however, the rule set may specify that the flow instead proceeds to a process stage **282**. Process stage **282** may generate a query relating to the billing alias that was presented to the cardholder. For example, the query may ask whether the cardholder is aware that the billing alias is used by that particular merchant. In some embodiments, process stage **282** may instead generate a message that simply informs the cardholder that the billing alias corresponds to the merchant, without posing a question.

The rule set may specify that process stage **284** generates one or more default queries. For example, one default query may ask whether the cardholder lent his or her card to a friend or family member around the time of the transaction. In some embodiments and/or scenarios, process stage **284** may be omitted from process flow **270**. Generally, the queries (and possibly non-query messages) generated in process flow **270** may serve to help the cardholder recall whether the transaction was made or authorized, and/or process flow **270** may prompt the cardholder for responses that are considered by others (e.g., personnel of an entity associated with FRMS **14** of FIG. 1) to determine whether the transaction was likely fraudulent.

Although not shown in FIG. 4E, in some embodiments process flow **270** may include a number of iterative stages in which responses are received from the cardholder (e.g., from the respective one of cardholder computing devices **20** in FIG. 1) and used to generate additional, more detailed questions for the cardholder. For example, if a first query asks whether the cardholder recalls personally making another “card present” transaction that occurred at a nearby time and place, and the cardholder responds “no,” a new query may be generated asking whether the cardholder

recalls personally making the next closest transaction (in terms of time and/or location).

F. Exemplary Document Fraud Detection Rule Set

Referring next to FIG. 4F, an exemplary rule set **290** (e.g., generated at process stage **204** of FIG. 3F) may use various factors relating to an imaged (e.g., photographed or scanned) physical document to determine whether the document should be flagged as a candidate for a more in-depth (e.g., manual) analysis/review for fraud purposes. The rule set **290** may correspond to a particular embodiment and scenario in which the document is one that includes at least a signature field (e.g., a personal check, a driver's license, etc.).

The factors considered under the rule set **290** may include a number of counterfeit factors **292** and a number of forgery factors **294**, each of which may be evaluated by image analysis unit **52** of FIG. 1 using one or more image processing techniques. The counterfeit factors **292** may relate to the look, presentation, format and/or structure of the document, while the forgery factors **294** may relate to the substance, style or format of information entered in one or more fields of the document.

As seen in FIG. 4F, the counterfeit factors **292** may include: (1) whether one or more absolute or relative dimensions and/or angles of the document, or of lines, illustrations, patterns, etc. shown on the document (excluding user-entered contents in fields such as the signature line), are outside one or more predetermined tolerances; (2) whether one or more colors on the document are outside a predetermined tolerance (e.g., color/frequency range); (3) whether one or more line thicknesses of the document (excluding user-entered field contents) are outside one or more predetermined tolerances; and/or (4) whether one or more fonts on the document (excluding user-entered field contents) are outside one or more predetermined tolerances. For example, image analysis unit **52** may determine whether the ratio of the document length to the document width is within 0.1% of an expected value. As another example, image analysis unit **52** may determine whether horizontal and vertical lines on the document are within 0.3 degrees of the horizontal and vertical edges of the document, respectively. As yet another example, image analysis unit **52** may determine whether a font used for a field descriptor or other text on the document matches an expected font (e.g., by meeting a similarity threshold measured in any suitable manner). In other embodiments, the counterfeit factors **292** may include more, fewer and/or different factors than those shown in FIG. 4F.

The forgery factors **294** may include: (1) whether a signature entered in a signature field of the document match is outside a predetermined tolerance (e.g., using any suitable signature recognition technique); (2) whether handwriting entered in one or more fields of the document is outside a predetermined tolerance (e.g., by applying a suitable handwriting recognition technique); and/or (3) whether the format of information entered by a user in one or more fields does not match an expected format (e.g., using "9.12.16" rather than the expected "9/12/2016," as established based upon other documents known to have been populated and/or submitted by the purported applicant). In other embodiments, the forgery factors **294** may include more, fewer and/or different factors than those shown in FIG. 4F.

Generally, the data indicative of whether the circumstances corresponding to counterfeit factors **292** and/or forgery factors **294** are present/true for a particular document may be included in the first document image data **210** described above in connection with FIG. 3F.

As is also seen in FIG. 4F, each of the counterfeit factors **292** and forgery factors **294** may be associated with a

particular score or weighting value. In the rule set **290** shown in FIG. 4F, a total score may be calculated based upon which factors are, or are not, present. In some embodiments, certain factors may instead be associated with negative scores, and/or certain factors may be associated with metrics or algorithms that determine how heavily those factors are weighed.

The rule set **290** may then output the total score, a normalized total score, an indication of whether the total score exceeded a threshold, a probability calculated based upon the total score, and/or some other indicator or measure of the likelihood that the document is fraudulent. Alternatively, the rule set **290** may output a separate total score, normalized score, probability, or other metric, for each of counterfeit factors **292** and forgery factors **294**, with the counterfeit metric indicating the likelihood that the document is a counterfeit and the forgery metric indicating the likelihood that the document was fraudulently populated by someone other than the purported person (e.g., by someone other than the person corresponding to the name, signature, address, etc. on the document). In the example shown in FIG. 4F, it can be seen that larger scores generally correspond to a greater probability that the document is fraudulent. In some embodiments, the rule set **290** also includes one or more other types of factors not shown in FIG. 4F, and/or omits either counterfeit factors **292** or forgery factors **294**.

V. Exemplary Methods for Fraud Detection & Classification

FIGS. 5-11 depict flow diagrams of various exemplary computer-implemented methods that may be implemented by one or more components of AFSS **12** of FIG. 1. In one embodiment, AFSS **12** implements all of the methods corresponding to FIGS. 5-11. In other embodiments, AFSS **12** implements only a subset (e.g., one, two, etc.) of the methods corresponding to FIGS. 5-11. Each of the methods described below may be implemented by fraud detection/classification unit **36** of FIG. 1, for example.

A. Exemplary Methods for Detecting Chargeback Candidates and/or Compromised Merchants

Referring now to FIG. 5, an exemplary computer-implemented method **300** may be used to identify a potential chargeback scenario, e.g., to narrow down the field of fraudulent transactions that should be further investigated for chargeback purposes. In the method **300**, chargeback candidate detection rules may be generated or updated (block **301**). The rules may be generated or updated using a machine learning program, such as any of the types of machine learning programs discussed above in connection with ML rule generator **40** of FIG. 1 or process stage **84** of FIG. 2, for example. The machine learning program may be trained using transaction data associated with a plurality of financial transactions, and chargeback determinations each corresponding to a respective one of the plurality of financial transactions. The chargeback determinations (e.g., indications of whether a chargeback was or was not instituted for each transaction) may each be made in accordance with chargeback rules associated with (e.g., provided by) a card network entity.

An indication that fraud has been confirmed for a first financial transaction, associated with a particular financial account and a particular merchant, may be received (block **302**). For example, an automated investigation (e.g., using any suitable process, method or technique described herein), and/or a manual investigation, may have been performed to determine that the first financial transaction was fraudulent in some way, and a user input or other data indicative of that outcome may be received at block **302**.

35

First transaction data, associated with the first financial transaction, may be retrieved (block 303). For example, a database containing account records (e.g., account records database 30 of FIG. 1) may contain a record of transactions associated with each of multiple accounts, and the transaction of interest may be retrieved at block 303 by accessing the information stored in the database.

It may be determined, by applying the chargeback candidate detection rules generated or updated at block 301 to the first transaction data retrieved at block 303, that a chargeback may be warranted for the first financial transaction (block 304). For example, the first financial transaction may be flagged as a relatively strong candidate with respect to the appropriateness of a chargeback payment from the merchant or acquiring bank to the card issuer. In other scenarios, not represented by FIG. 5, it may instead be determined that a chargeback is not a good chargeback candidate.

The determination at block 304 may be based upon different factors, in accordance with the specific scenario and/or embodiment. For example, the determination may be based at least upon a card associated with the financial account being expired when the first financial transaction occurred. As other examples, the determination may be based at least upon a dollar amount of the first financial transaction (e.g., the amount exceeding a threshold), and/or upon a second financial transaction under the same account duplicating one or more characteristics of the first financial transaction (e.g., the same amount, date, time, and/or other characteristics that may be indicative of an improper duplicate charge).

An indication that the chargeback may be warranted may be caused to be displayed to one or more people via one or more respective computing device user interfaces (block 305). The indication may also specify the transaction at issue, the merchant, and possibly other information such as the transaction date, amount, etc. The indication may be sent to a computing device of a card issuer or other entity, for example, to prompt a full review of the first financial transaction under the chargeback rules of the card network entity. Block 305 may be implemented by notification unit 56 of FIG. 1, for example.

In some embodiments, the method 300 may include one or more additional blocks not shown in FIG. 5. For example, the method 300 may include an additional block in which it is determined, by applying the chargeback rules of the card network entity to the first transaction data and additional data associated with that transaction (e.g., additional data received from the merchant and/or other sources), that a chargeback is indeed warranted for the transaction. The additional data may include data indicative of a card information entry mode used to conduct the first financial transaction (e.g., "card present," swiped, inserted in an EMV chip card reader, etc.), for example. In other embodiments and/or scenarios, the chargeback rules may be applied manually.

As another example, the method 300 may include an additional block in which first account data associated with the first financial account is retrieved (e.g., card expiration date, whether the card is an EMV card, etc., possibly obtained from the same account records database accessed at block 303). In such an embodiment, block 304 may further include applying the chargeback candidate detection rules to the first account data.

FIG. 6 illustrates an exemplary computer-implemented method 310 of determining that a merchant computer terminal warrants a chargeback after a consumer dispute or fraudulent transaction based upon dispute rules (or charge-

36

back rules). The method 310 may include, via one or more processors and/or transceivers, determining a merchant and/or a merchant computing terminal associated with a fraudulent financial transaction, such as by IP address (block 311).

The method 310 may include, the one or more processors and/or transceivers, retrieving a configuration, or parameters or specifications of the merchant computing terminal from a memory unit, or receiving such information via wireless communication or data transmission over one or more radio links or wireless communication channel transmitted by the merchant computing terminal (block 312). The method 310 may include retrieving, via the one or more processors and/or transceivers, current or up-to-date dispute rules regarding computer specifications or requirements (block 313). The method 310 may include comparing, via the one or more processors, the current dispute rules regarding computer requirements with the merchant computer terminal specifications to determine whether or not the merchant computer terminal is compliant with current dispute rules (block 314). For instance, the one or more processors may analyze the dispute rules to identify requirements or required specifications of merchant computing terminals. The method 310 may include verifying other dispute rules are satisfied for a valid chargeback and/or whether the financial transaction is still within requisite time or allowable recovery period. The method 310 may include, if the merchant computer terminal is not compliant with the dispute rules or computer required specifications, generating an electronic chargeback notification and transmitting it to a merchant computing device (block 315) to facilitate merchant notification.

For instance, the dispute rules may require merchant computing terminals to be equipped with "chip" readers, smart card readers, or the like. The method 310 may determine that the merchant computing terminal is only equipped with older or "swipe" technology, and if so, a card issuer may be entitled to a chargeback associated with a fraudulent transaction from the merchant.

In one aspect, a computer-implemented method of identifying a merchant computer terminal warranting a chargeback may be provided. The method may include (1) identifying, via one or more processors and/or transceivers, a merchant computer terminal (and/or an IP address of the merchant computer terminal) associated with a fraudulent financial transaction; (2) retrieving or receiving, via the one or more processors and/or transceivers, an actual configuration, parameters, and/or specifications associated with the merchant computer terminal; (3) retrieving or receiving, via the one or more processors and/or transceivers, up-to-date dispute rules associated with chargebacks, the dispute rules including merchant computer terminal requirements; (4) analyzing, via the one or more processors, the dispute rules to determine merchant computer terminal requirements; (5) comparing, via the one or more processors and/or transceivers, the merchant computer terminal requirements from the current dispute rules with the actual configuration, parameters, and/or specifications associated with the merchant computer terminal to identify if the merchant computer terminal is non-complaint with the dispute rules; (6) if the merchant computer terminal is non-complaint, then generating, via the one or more processors, an electronic notification that includes an identification of the financial transaction at issue, a transaction amount, and that a chargeback is warranted due to a merchant computing device being non-compliant with the dispute rules; and/or (7) transmitting, via the one or more processors and/or transceivers, the electronic notification to a merchant computing device via

37

wireless communication over one or more radio frequency links to facilitating resolving financial transaction disputes.

The method may include receiving, via the one or more processors and/or transceivers, dispute rules associated with a credit card issuer; and identifying, via the one or more processors, merchant computer terminal technical specifications that would warrant a chargeback and merchant computer terminal technical specifications that would not warrant a chargeback.

The method may include receiving, via the one or more processors and/or transceivers, actual technical specifications of a merchant computer associated with the financial transaction; comparing, via the one or more processors, the actual technical specifications with the merchant computer terminal technical specifications that would warrant a chargeback to determine that a chargeback is warranted; and/or generating, via the one or more processors, an electronic notification indicating that the chargeback is warranted.

The method may include receiving, via the one or more processors and/or transceivers, actual technical specifications of a merchant computer associated with the financial transaction; comparing, via the one or more processors, the actual technical specifications with the merchant computer terminal technical specifications that would not warrant a chargeback (or satisfy computer requirements of the dispute rules) to determine that a chargeback is not warranted; and/or generating, via the one or more processors, an electronic notification indicating that the chargeback is not warranted based upon computer specifications.

In another aspect, a computer system configured to identify a merchant computer terminal warranting a chargeback may be provided. The computer system may include one or more processors and/or transceivers configured to: identify a merchant computer terminal (and/or an IP address of the merchant computer terminal) associated with a fraudulent financial transaction; retrieve from a local or remote memory unit, and/or receive via wireless communication or data transmission over one or more radio links or wireless communication channels, an actual configuration, parameters, and/or specifications associated with the merchant computer terminal; retrieve from a local or remote memory unit, and/or receive via wireless communication or data transmission over one or more radio links or wireless communication channels, up-to-date dispute rules associated with chargebacks, the dispute rules including merchant computer terminal requirements; determine merchant computer terminal requirements from processor analysis of the dispute rules; compare the merchant computer terminal requirements from the current dispute rules with the actual configuration, parameters, and/or specifications associated with the merchant computer terminal to identify if the merchant computer terminal is non-complaint with the dispute rules; if the merchant computer terminal is non-complaint, then generate an electronic notification that includes an identification of the financial transaction at issue, a transaction amount, and that a chargeback is warranted due to a merchant computing device being non-compliant with the dispute rules; and/or transmit the electronic notification to a merchant computing device via wireless communication or data transmission over one or more radio frequency links or wireless communication links to facilitating resolving financial transaction disputes.

FIG. 7 illustrates a computer-implemented method 320 of identifying compromised, or potentially compromised, merchants and/or reducing future chargebacks (or actual losses) or fraud. The method 320 may include, via one or more

38

processors and/or transceivers (such as via wireless communication or data transmission over one or more radio frequency links and/or wireless communication channels), (1) generating a list of chargebacks for a given period of time (block 321); (2) inputting the list of chargebacks associated with the merchant (and/or associated facts related to the underlying financial transaction) into a machine learning program that is trained to identify a reason why each chargeback was generated (or actual loss was incurred), and/or a type or classification of each chargeback (block 322) (such as stolen card, lost card, account take over, fraudulent application, compromised computing system, identity theft, counterfeit card, etc.); (3) analyzing the reason and/or the type or classification of each chargeback to identify those chargebacks associated with a potentially compromised merchant computer system (block 323) (such as by inputting the reason and/or the type or classification of each chargeback into a second machine learning program trained to identify chargebacks associated with potentially compromised merchant computer systems); (4) determining a merchant associated with a number of chargebacks within the given period of time (including those that are identified as being associated with potentially compromised merchant computer systems) that are over a predetermined threshold amount or an outlier amount indicative of suspicious activity (block 324), such as a spike in a specific type or classification of chargeback that is an individual merchant is experiencing; (5) generating an electronic notification notifying the potentially compromised merchant of a possibility of a data breach or other security concern (block 325); and/or (6) transmitting the electronic notification to a compromised merchant computing device (block 326) to alert the compromised merchant of the data breach or other security concern.

FIG. 8 illustrates a computer-implemented method 330 of identifying compromised, or potentially compromised, merchants. The method 330 may include, via one or more processors and/or transceivers (such as via wireless communication or data transmission over one or more radio frequency links and/or wireless communication channels), (1) generating financial fraud electronic alerts using a rules-based model or engine (block 331); (2) inputting the financial fraud electronic alerts into a machine learning program trained to determine a reason why each fraud alert was generated (such as stolen card, lost card, account take over, fraudulent application, compromised computing system, identity theft, counterfeit card, etc.), and/or a type or classification of each fraud alert (block 332); (3) analyzing the reason and/or the type or classification of each fraud alert to identify those fraud alerts associated with a potentially compromised merchant computer system (block 333) (such as by inputting the reason and/or the type or classification of each fraud alert into a second machine learning program trained to identify fraud alerts associated with potentially compromised merchant computer systems or determine unusual fraud alerts for individual merchants or types of merchants); (4) determining a merchant associated with a number of fraud alerts (such as those identified as being associated with potentially compromised merchant computer systems, or other types of alerts) that are over a predetermined threshold or an outlier amount indicative of suspicious activity (block 334); (5) generating an electronic notification notifying the compromised merchant of a possibility of a data breach or other security concern (block 335); and/or (6) transmitting the electronic notification to a compromised merchant computing device (block 336) to

alert the compromised, or potentially compromised, merchant of the data breach or other security concern.

H. Exemplary Application Fraud Detection Methods

Referring now to FIG. 9, an exemplary computer-implemented method **340** may be used to detect potential application fraud, e.g., to detect instances where an applicant for a loan, credit, checking, savings, investment and/or other type of account, service or benefit is not the person that he or she purports to be. In the method **340**, application data, indicative of applicant entries in one or more fields of an application associated with a potential account, may be received (block **341**). The application may be an application received from a source computing device via the Internet (e.g., via a web page, via a dedicated application, via email, etc.), for example.

The source computing device, via which the applicant submitted the application data, may be identified (block **342**). The source computing device may be similar to one of computing devices **20** of FIG. 1, for example. In one embodiment, block **342** may include identifying an IP address of the source computing device.

Search history data, indicative of one or more search terms submitted to an Internet-based search engine (e.g., a Google® or Yahoo® search engine) via the source computing device, may be retrieved (block **343**). The search history data may be retrieved from a database maintained by an ISP of the applicant, or a temporary database that is locally constructed based upon information received directly from the source computing device, for example. In one embodiment where the IP address of the source computing device is identified at block **342**, the search history data may include search terms that were submitted to the Internet-based search engine from the identified IP address.

It may be determined, by analyzing the application data and the search history data, that the applicant may not be the person whom the applicant purports to be (block **344**), e.g., the person whose name and/or signature is included on the application, or is otherwise listed or indicated in connection with the application. For example, the application may be flagged at block **344** for a further, more in-depth review. In other scenarios, not represented by FIG. 9, it may instead be determined that the applicant likely is the person whom he or she purports to be, such that it is not necessary to flag the application for further review.

Block **344** may include various operations, depending upon the embodiment and/or scenario. For example, block **344** may include comparing information included in retrieved search term(s) (and/or included in search results) to information included in the applicant entries in the application fields. As other examples, block **344** may include comparing a name included in the search term(s) to a name included in the applicant entries, determining that the search term(s) is/are directed to discovering an address associated with the name (e.g., by including the terms “Jonathon Doe address” or “Jonathon Doe residence,” etc.), and/or determining that the search term(s) is/are directed to discovering an employment history associated with the name (e.g., by including the terms “Jonathon Doe employer” or “Jonathon Doe work history,” etc.).

An indication of potential application fraud may be caused to be displayed to one or more people via one or more respective computing device user interfaces (block **345**). The indication may also specify the purported applicant, the actual applicant, and/or other relevant information. The indication may be sent to a computing device of a card issuer or other entity, for example, to prompt a full review of the

application (e.g., a manual investigation). Block **345** may be implemented by notification unit **56** of FIG. 1, for example.

In some embodiments, the method **340** may include one or more additional blocks not shown in FIG. 9. For example, in an embodiment where block **344** includes applying application fraud detection rules to the application data and the search history data, the method **340** may include an additional block in which the application fraud detection rules are generated or updated. For example, the rules may be generated or updated by training a machine learning program, such as any of the types of machine learning programs discussed above in connection with ML rule generator **40** of FIG. 1 or process stage **84** of FIG. 2. The machine learning program may be trained using historical application fraud determinations made in connection with a plurality of other applications, and additional search history data associated with additional source computing devices (e.g., the devices that submitted the data for those applications), for example.

FIG. 10 illustrates an exemplary computer-implemented method **350** of using IP address and/or browsing activity to identify fraudulent online or virtual applications. The method **350** may include, via one or more processors and/or transceivers, receiving an online or virtual application via wireless communication or data transmission over one or more radio links or wireless communication channels (block **351**), such as an application for financial services, such as vehicle, home, or personal loans, or credit cards. The virtual application may include an IP address identification and/or a recent browser activity history of the computer. The method **350** may include determining or retrieving a name of the applicant listed on the online application (block **352**). The method **350** may include determining an IP address of the source computer that generated the online application (block **353**). The method **350** may include retrieving or receiving recent online browsing activity of the source computer (such as by using the IP address and/or browsing activity received with the virtual application, or querying the source computer for such information) (block **354**). The method **350** may include determining if the recent online activity of the source computer includes internet searches on the applicant's name and/or the applicant's social media presence or website (block **355**) (indicating that the person submitting the application retrieved facts about the named applicant in an effort to perform fraudulent activity or impersonate the named applicant). If so, the method **350** may include flagging the virtual application as fraudulent, and/or in need of further manual or processor review (block **356**).

In one aspect, a computer-implemented method of using browsing activity to identify fraudulent online or virtual applications may be provided. The method may include (1) receiving, via one or more processors and/or transceivers, a virtual application over one or more radio frequency links or wireless communication channels; (2) determining, via the one or more processors, an applicant name on the virtual application; (3) determining, via the one or more processors, an IP address of a source computer from which the virtual application originated; (4) determining or receiving, via the one or more processors and/or transceivers, an online browsing or search history associated with the IP address and/or source computer; (5) determining, via the one or more processors, if the online browsing or search history indicates recent internet searches for or on the applicant name; and/or (6) if so, flagging, via the one or more processors, the virtual application as fraudulent and generating an electronic alert

41

indicating that the virtual application is fraudulent to facilitating identifying fraudulent virtual or online applications for goods or services.

The virtual application may include a field or routine that records and retrieves recent online browsing or searching activity of the source computer, and that information is transmitted from the source computer along with the virtual application when the virtual application is submitted. In other words the virtual application may carry with it online browsing or search activity associated with the source computer when it is transmitted or submitted to a remote server, such as a financial services provider remote server.

The virtual application may include a field or routine that records and retrieves recent social media searching activity of the source computer. The virtual application may include a field or routine that records and retrieves recent internet search engine searching activity of the source computer. The virtual application may include an IP address field of the source computer. The virtual application may include a current GPS location of the source computer.

The method may further include determining, via the one or more processors, whether the named applicant is an existing customer; if so, receiving, via the one or more processors, customer data from one or more sources indicating a current customer location; determining, via the one or more processors, whether the current customer location corresponds to, or is within a predetermined distance of, a location of the source computer from which the virtual application originated; and/or if the current customer location does not correspond to the location of the source computer, then generating, via the one or more processors, an electronic alert indicating such.

Determining an online browsing or search history associated with the IP address and/or source computer may include retrieving an online browsing or search history stored with or within the virtual application. Receiving an online browsing or search history associated with the IP address and/or source computer may include receiving, via the one or more processors and/or transceivers, an online browsing or search history over one or more radio links or wireless communication channels via wireless communication or data transmission. The virtual or online application may be an application for financial services or products, such vehicle, home, or personal loans; a credit card application; or an application for auto, homeowners, or life insurance.

In another aspect, a computer system configured to use IP addresses and/or browsing activity to identify fraudulent online or virtual applications may be provided. The computer system may include one or more processors and/or transceiver configured to: receive a virtual application via wireless communication or data transmission over one or more radio frequency links or wireless communication channels; determine an applicant name on the virtual application (or retrieve the virtual applicant name saved within the virtual application); determine an IP address of a source computer from which the virtual application originated; determine or receive an online browsing or search history associated with the IP address and/or source computer; determine if the online browsing or search history indicates recent internet searches for, or on, the applicant name; and/or if so, flag the virtual application as fraudulent and generate an electronic alert indicating that the virtual application is fraudulent to facilitating identifying fraudulent virtual or online applications for goods or services.

FIG. 11 illustrates a computer-implemented method 360 of identifying mortgage fraud. The method 360 may include,

42

via one or more processors and/or transceivers (such as via wireless communication or data transmission over one or more radio frequency links or wireless communication channels), (1) receiving a virtual application for a home loan or line of credit for a property and an identity of a person or entity applying (block 361); (2) receiving home loan or line of credit information for the property from multiple sources or lenders, each of the multiple sources or lenders having provided a previous home loan or line of credit for the property (block 362), such as providing a previous home loan or line of credit for the property to the person or entity applying for another home loan or line of credit; (3) retrieving from a memory, or receiving an appraisal for the property, the appraisal including an estimated or actual value of the property (block 363), such as receiving a virtual appraisal from another lender; (4) determining that a total amount of the multiple home loans or lines of credit for the property exceed the estimated or actual value of the property, or a percentage thereof (block 364); (5) if so, generating an electronic notification indicating that mortgage fraud associated with the property may exist and/or denying the application (block 365); and/or (6) transmitting the electronic notification to the multiple sources or lenders to notify them of the potential mortgage fraud (block 366).

VI. Exemplary System for Fraud Detection & Classification

FIG. 12 depicts an exemplary computer system 500 in which the techniques described herein may be implemented, according to one embodiment. The computer system 500 of FIG. 12 may include a computing device in the form of a computer 510. Components of the computer 510 may include, but are not limited to, a processing unit 520, a system memory 530, and a system bus 521 that couples various system components including the system memory 530 to the processing unit 520. The system bus 521 may be any of several types of bus structures including a memory bus or memory controller, a peripheral bus, or a local bus, and may use any suitable bus architecture. By way of example, and not limitation, such architectures include the Industry Standard Architecture (ISA) bus, Micro Channel Architecture (MCA) bus, Enhanced ISA (EISA) bus, Video Electronics Standards Association (VESA) local bus, and Peripheral Component Interconnect (PCI) bus (also known as Mezzanine bus).

Computer 510 may include a variety of computer-readable media. Computer-readable media may be any available media that can be accessed by computer 510 and may include both volatile and nonvolatile media, and both removable and non-removable media. By way of example, and not limitation, computer-readable media may comprise computer storage media and communication media. Computer storage media may include volatile and nonvolatile, removable and non-removable media implemented in any method or technology for storage of information such as computer-readable instructions, data structures, program modules or other data. Computer storage media may include, but is not limited to, RAM, ROM, EEPROM, FLASH memory or other memory technology, CD-ROM, digital versatile disks (DVD) or other optical disk storage, magnetic cassettes, magnetic tape, magnetic disk storage or other magnetic storage devices, or any other medium which can be used to store the desired information and which can be accessed by computer 510.

Communication media typically embodies computer-readable instructions, data structures, program modules or other data in a modulated data signal such as a carrier wave or other transport mechanism, and may include any information delivery media. The term "modulated data signal"

means a signal that has one or more of its characteristics set or changed in such a manner as to encode information in the signal. By way of example, and not limitation, communication media may include wired media such as a wired network or direct-wired connection, and wireless media such as acoustic, radio frequency (RF), infrared and other wireless media. Combinations of any of the above are also included within the scope of computer-readable media.

The system memory 530 may include computer storage media in the form of volatile and/or nonvolatile memory such as read only memory (ROM) 531 and random access memory (RAM) 532. A basic input/output system 533 (BIOS), containing the basic routines that help to transfer information between elements within computer 510, such as during start-up, is typically stored in ROM 531. RAM 532 typically contains data and/or program modules that are immediately accessible to, and/or presently being operated on, by processing unit 520. By way of example, and not limitation, FIG. 12 illustrates operating system 534, application programs 535, other program modules 536, and program data 537.

The computer 510 may also include other removable/non-removable, volatile/nonvolatile computer storage media. By way of example only, FIG. 12 illustrates a hard disk drive 541 that reads from or writes to non-removable, nonvolatile magnetic media, a magnetic disk drive 551 that reads from or writes to a removable, nonvolatile magnetic disk 552, and an optical disk drive 555 that reads from or writes to a removable, nonvolatile optical disk 556 such as a CD ROM or other optical media. Other removable/non-removable, volatile/nonvolatile computer storage media that can be used in the exemplary operating environment include, but are not limited to, magnetic tape cassettes, flash memory cards, digital versatile disks, digital video tape, solid state RAM, solid state ROM, and the like. The hard disk drive 541 may be connected to the system bus 521 through a non-removable memory interface such as interface 540, and magnetic disk drive 551 and optical disk drive 555 may be connected to the system bus 521 by a removable memory interface, such as interface 550.

The drives and their associated computer storage media discussed above and illustrated in FIG. 12 provide storage of computer-readable instructions, data structures, program modules and other data for the computer 510. In FIG. 12, for example, hard disk drive 541 is illustrated as storing operating system 544, application programs 545, other program modules 546, and program data 547. Note that these components can either be the same as or different from operating system 534, application programs 535, other program modules 536, and program data 537. Operating system 544, application programs 545, other program modules 546, and program data 547 are given different numbers here to illustrate that, at a minimum, they are different copies. A user may enter commands and information into the computer 510 through input devices such as cursor control device 561 (e.g., a mouse, trackball, touch pad, etc.) and keyboard 562. A monitor 591 or other type of display device is also connected to the system bus 521 via an interface, such as a video interface 590. In addition to the monitor, computers may also include other peripheral output devices such as printer 596, which may be connected through an output peripheral interface 595.

The computer 510 may operate in a networked environment using logical connections to one or more remote computers, such as a remote computer 580. The remote computer 580 may be a personal computer, a server, a router, a network PC, a peer device or other common network node,

and may include many or all of the elements described above relative to the computer 510, although only a memory storage device 581 has been illustrated in FIG. 12. The logical connections depicted in FIG. 12 include a local area network (LAN) 571 and a wide area network (WAN) 573, but may also include other networks. Such networking environments are commonplace in hospitals, offices, enterprise-wide computer networks, intranets and the Internet.

When used in a LAN networking environment, the computer 510 is connected to the LAN 571 through a network interface or adapter 570. When used in a WAN networking environment, the computer 510 may include a modem 572 or other means for establishing communications over the WAN 573, such as the Internet. The modem 572, which may be internal or external, may be connected to the system bus 521 via the input interface 560, or other appropriate mechanism. The communications connections 570, 572, which allow the device to communicate with other devices, are an example of communication media, as discussed above. In a networked environment, program modules depicted relative to the computer 510, or portions thereof, may be stored in the remote memory storage device 581. By way of example, and not limitation, FIG. 12 illustrates remote application programs 585 as residing on memory device 581.

The techniques for detecting and/or classifying fraud described above may be implemented in part or in their entirety within a computer system such as the computer system 500 illustrated in FIG. 12. The computer 510 may be included in AFSS 12 of FIG. 1, for example, and/or the remote application programs 585 may include one or more applications of either FAMS 14, one of cardholder computing device 20, one of merchant computing systems 22, or a computing device of other sources 24. Moreover, the functionality of fraud detection/classification unit 36 of FIG. 1 may be implemented by one or more of application programs 535 and/or other program modules 536. As another example, ML rules database 58, account holder behaviors database 60 and/or chargeback rules database 62 of FIG. 1 may be stored in hard disk drive 541 (e.g., as program data 547), magnetic disk 552 and/or optical disk drive 555, and/or the data retrieved by fraud detection/classification unit 36 of FIG. 1 may be stored in hard disk drive 541 (e.g., as program data 547) and/or RAM 532 (e.g., as program data 537).

VII. Exemplary Method Embodiments

In another aspect, a computer-implemented method, implemented in one or more servers or other computing devices, of identifying a potential chargeback scenario may include (1) generating or updating, by one or more processors of the one or more servers, chargeback candidate detection rules, at least by training a machine learning program using at least (i) transaction data associated with a plurality of financial transactions, and (ii) chargeback determinations each made in accordance with chargeback rules associated with a card network entity, and each corresponding to a respective one of the plurality of financial transactions; (2) receiving, by the one or more processors, an indication that fraud has been confirmed for a first financial transaction associated with a merchant and a first financial account; (3) retrieving, by the one or more processors and from an account records database, first transaction data associated with the first financial transaction; (4) determining, by the one or more processors applying the chargeback candidate detection rules to the first transaction data, that a chargeback may be warranted for the first financial transaction; and/or (5) causing, by the one or more processors, an indication that the chargeback may be warranted to be

45

displayed to one or more people via one or more respective computing device user interfaces, wherein the indication may further specifying at least the first financial transaction and the merchant. The method may include additional, fewer or alternative actions, such as any of those discussed elsewhere herein.

For instance, the method may further include determining, by the one or more processors applying the chargeback rules to (i) the first transaction data and (ii) additional data associated with the first financial transaction, that a chargeback is warranted for the first financial transaction.

Additionally or alternatively, the method may further include, prior to determining that the chargeback is warranted, receiving at least a portion of the additional data from the merchant. Additionally or alternatively, the portion of the additional data received from the merchant may include data indicative of a card information entry mode used to conduct the first financial transaction.

Additionally or alternatively, the method may further include retrieving, by the one or more processors and from the account records database, first account data associated with the first financial account, and determining that the chargeback may be warranted for the first financial transaction may further be performed by the one or more processors applying the chargeback candidate detection rules to the first account data.

Additionally or alternatively, determining that a chargeback may be warranted for the first financial transaction may include determining that a chargeback may be warranted based at least upon a card associated with the first financial account being expired when the first financial transaction occurred. Additionally or alternatively, determining that a chargeback may be warranted for the first financial transaction may include determining that a chargeback may be warranted based at least upon (i) a dollar amount of the first financial transaction, and/or (ii) a second financial transaction duplicating one or more characteristics of the first financial transaction.

In another aspect, a computer-implemented method, implemented in one or more servers or other computing devices, of detecting potential application fraud may include (1) receiving, by one or more processors of the one or more servers, application data indicative of applicant entries in one or more fields of an application associated with a potential account; (2) identifying, by the one or more processors, a source computing device via which the applicant submitted the application data; (3) retrieving, by the one or more processors and from a database, search history data indicative of one or more search terms submitted to an Internet-based search engine via the source computing device; (4) determining, by the one or more processors analyzing the application data and the search history data, that the applicant may not be a person whom the applicant purports to be; and/or (5) causing, by the one or more processors, an indication of potential application fraud to be displayed to one or more people via one or more respective computing device user interfaces. The method may include additional, fewer or alternative actions, such as any of those discussed elsewhere herein.

For instance, identifying the source computing device may include identifying an Internet Protocol (IP) address of the source computing device. Additionally or alternatively, retrieving search history data may include retrieving search history data indicative of search terms submitted to the Internet-based search engine from the identified IP address.

Additionally or alternatively, determining that the applicant may not be the person whom the applicant purports to

46

be may include comparing information included in the one or more search terms to information included in the applicant entries in the one or more fields; comparing a name included in the one or more search terms to a name included in the applicant entries in the one or more fields; determining that the applicant may not be the person whom the applicant purports to be may further include determining that the one or more search terms are directed to discovering an address associated with the name; and/or determining that the one or more search terms are directed to discovering an employment history associated with the name.

Additionally or alternatively, (1) determining that the applicant is not, or may not be, the person whom the applicant purports to be may include determining, by the one or more processors applying application fraud detection rules to the application data and the search history data, that the applicant may not be a person whom the applicant purports to be; and/or (2) the method may further include generating or updating, by the one or more processors, the application fraud detection rules, at least by training a machine learning program using at least (i) historical application fraud determinations made in connection with a plurality of applications, and (ii) additional search history data associated with a plurality of additional source computing devices.

VIII. Exemplary System Embodiments

In another aspect, a computer system for identifying a potential chargeback scenario may include (1) an account records database configured to store data associated with a plurality of financial accounts; (2) a rules database configured to store chargeback candidate detection rules; (3) one or more processors; and/or (4) a non-transitory memory. The non-transitory memory stores instructions that, when executed by the one or more processors, may cause the one or more processors to (1) generate or update the chargeback candidate detection rules, at least by training a machine learning program using at least (i) transaction data associated with a plurality of financial transactions stored in the account records database, and (ii) chargeback determinations each made in accordance with chargeback rules associated with a card network entity, and each corresponding to a respective one of the plurality of financial transactions; (2) receive an indication that fraud has been confirmed for a first financial transaction associated with a merchant and a first financial account; (3) retrieve, from the account records database, first transaction data associated with the first financial transaction; (4) determine, by applying the chargeback candidate detection rules stored in the rules database to the first transaction data, that a chargeback may be warranted for the first financial transaction; and/or (5) cause an indication that the chargeback may be warranted to be displayed to one or more people via one or more respective computing device user interfaces, wherein the indication may further specify at least the first financial transaction and the merchant. The system may include additional, fewer or alternative components, features and/or functionality, such as any of those discussed elsewhere herein.

For instance, the instructions may further cause the one or more processors to determine, by applying the chargeback rules to (i) the first transaction data and (ii) additional data associated with the first financial transaction, that a chargeback is warranted for the first financial transaction. Additionally or alternatively, the instructions may further cause the one or more processors to, prior to determining that the chargeback is warranted, receive at least a portion of the additional data from the merchant.

Additionally or alternatively, the portion of the additional data received from the merchant may include data indicative of a card information entry mode used to conduct the first financial transaction. Additionally or alternatively, the instructions may further cause the one or more processors to retrieve, from the account records database, first account data associated with the first financial account, and/or the instructions may cause the one or more processors to determine that the chargeback may be warranted for the first financial transaction further by applying the chargeback candidate detection rules to the first account data.

Additionally or alternatively, the instructions may cause the one or more processors to determine that a chargeback may be warranted for the first financial transaction based at least upon a card associated with the first financial account being expired when the first financial transaction occurred. Additionally or alternatively, the instructions may cause the one or more processors to determine that a chargeback may be warranted for the first financial transaction based at least upon (i) a dollar amount of the first financial transaction, and/or (ii) a second financial transaction duplicating one or more characteristics of the first financial transaction.

In another aspect, a computer system for detecting potential application fraud may include (1) a search history database configured to store data associated with a plurality of Internet searches; (2) one or more processors; and/or (3) a non-transitory memory. The non-transitory memory stores instructions that, when executed by the one or more processors, may cause the one or more processors to (1) receive application data indicative of applicant entries in one or more fields of an application associated with a potential account; (2) identify a source computing device via which the applicant submitted the application data; (3) retrieve, from the search history database, search history data indicative of one or more search terms submitted to an Internet-based search engine via the source computing device; (4) determine, by analyzing the application data and the search history data, that the applicant may not be a person whom the applicant purports to be; and/or (5) cause an indication of potential application fraud to be displayed to one or more people via one or more respective computing device user interfaces. The system may include additional, fewer or alternative components, features and/or functionality, such as any of those discussed elsewhere herein.

For instance, the instructions may cause the one or more processors to identify the source computing device at least by identifying an Internet Protocol (IP) address of the source computing device. Additionally or alternatively, the search history data may be indicative of search terms submitted to the Internet-based search engine from the identified IP address.

The instructions may cause the one or more processors to determine that the applicant may not be the person whom the applicant purports to be (i) at least by comparing information included in the one or more search terms to information included in the applicant entries in the one or more fields; (ii) at least by comparing a name included in the one or more search terms to a name included in the applicant entries in the one or more fields; (iii) further by determining that the one or more search terms are directed to discovering an address associated with the name; and/or (iv) further by determining that the one or more search terms are directed to discovering an employment history associated with the name.

Additionally or alternatively, (1) the instructions may cause the one or more processors to determine that the applicant may not be the person whom the applicant pur-

ports to be at least by determining, by applying application fraud detection rules to the application data and the search history data, that the applicant may not be the person whom the applicant purports to be; and/or (2) the instructions may further cause the one or more processors to generate or update the application fraud detection rules, at least by training a machine learning program using at least (i) historical application fraud determinations made in connection with a plurality of applications, and (ii) additional search history data associated with a plurality of additional source computing devices.

IX. Exemplary Computer-Readable Medium Embodiments

In one aspect, a non-transitory, computer-readable medium storing instructions that, when executed by one or more processors, may cause the one or more processors to (1) generate or update chargeback candidate detection rules, at least by training a machine learning program using at least (i) transaction data associated with a plurality of financial transactions, and (ii) chargeback determinations each made in accordance with chargeback rules associated with a card network entity, and each corresponding to a respective one of the plurality of financial transactions; (2) receive an indication that fraud has been confirmed for a first financial transaction associated with a merchant and a first financial account; (3) retrieve, from an account records database, first transaction data associated with the first financial transaction; (4) determine, by applying the chargeback candidate detection rules to the first transaction data, that a chargeback may be warranted for the first financial transaction; and/or (5) cause an indication that the chargeback may be warranted to be displayed to one or more people via one or more respective computing device user interfaces, wherein the indication may further specify at least the first financial transaction and the merchant. The system may include additional, fewer or alternative components, features and/or functionality, such as any of those discussed elsewhere herein.

For instance, the instructions may further cause the one or more processors to determine, by applying the chargeback rules to (i) the first transaction data and (ii) additional data associated with the first financial transaction, that a chargeback is warranted for the first financial transaction. Additionally or alternatively, the instructions may further cause the one or more processors to, prior to determining that the chargeback is warranted, receive at least a portion of the additional data from the merchant.

Additionally or alternatively, the portion of the additional data received from the merchant may include data indicative of a card information entry mode used to conduct the first financial transaction. Additionally or alternatively, the instructions may further cause the one or more processors to retrieve, from the account records database, first account data associated with the first financial account, and/or the instructions may cause the one or more processors to determine that the chargeback may be warranted for the first financial transaction further by applying the chargeback candidate detection rules to the first account data.

Additionally or alternatively, the instructions may cause the one or more processors to determine that a chargeback may be warranted for the first financial transaction based at least upon a card associated with the first financial account being expired when the first financial transaction occurred.

In another aspect, a non-transitory, computer-readable medium stores instructions that, when executed by one or more processors, may cause the one or more processors to (1) receive application data indicative of applicant entries in one or more fields of an application associated with a

49

potential account; (2) identify a source computing device via which the applicant submitted the application data; (3) retrieve, from a search history database, search history data indicative of one or more search terms submitted to an Internet-based search engine via the source computing device; (4) determine, by analyzing the application data and the search history data, that the applicant may not be a person whom the applicant purports to be; and/or (5) cause an indication of potential application fraud to be displayed to one or more people via one or more respective computing device user interfaces. The computer-readable medium may store instructions that include additional, fewer or alternative actions, such as any of those discussed elsewhere herein.

For instance, the instructions may cause the one or more processors to identify the source computing device at least by identifying an Internet Protocol (IP) address of the source computing device. Additionally or alternatively, the search history data may be indicative of search terms submitted to the Internet-based search engine from the identified IP address. Additionally or alternatively, the instructions may cause the one or more processors to determine that the applicant may not be the person whom the applicant purports to be at least by comparing information included in the one or more search terms to information included in the applicant entries in the one or more fields.

X. Additional Considerations

The following additional considerations apply to the foregoing discussion. Throughout this specification, plural instances may implement operations or structures described as a single instance. Although individual operations of one or more methods are illustrated and described as separate operations, one or more of the individual operations may be performed concurrently, and nothing requires that the operations be performed in the order illustrated. These and other variations, modifications, additions, and improvements fall within the scope of the subject matter herein.

What is claimed:

1. A computer-implemented method of using browsing activity to identify fraudulent online or virtual applications, the method comprising:

determining, by the one or more processors, an applicant name on a virtual application;

determining, by the one or more processors, an IP address of a source computer from which the virtual application originated;

determining, by the one or more processors, an online browsing or search history associated with the IP address;

determining, by the one or more processors, whether the online browsing or search history indicates recent Internet searches for the applicant name; and

in response to determining that the online browsing or search history does indicate recent Internet searches for the applicant name, (i) flagging, by the one or more processors, the virtual application as fraudulent and (ii) generating an electronic alert indicating that the virtual application is fraudulent.

2. The computer-implemented method of claim 1, wherein the virtual application includes a field or routine that records recent online browsing or searching activity of the source computer, and the method further comprises receiving a recorded recent online browsing or searching activity of the source computer along with the virtual application when the virtual application is submitted.

50

3. The computer-implemented method of claim 1, wherein the virtual application includes a field or routine that records recent social media searching activity of the source computer.

4. The computer-implemented method of claim 1, wherein the virtual application includes a field or routine that records recent Internet search engine searching activity of the source computer.

5. The computer-implemented method of claim 1, wherein the virtual application includes an IP address field of the source computer, and wherein determining the IP address of the source computer includes reading the IP address field of the virtual application.

6. The computer-implemented method of claim 1, wherein the virtual application includes a field indicating a current GPS location of the source computer.

7. The computer-implemented method of claim 1, the method further comprising:

determining, by the one or more processors, whether the applicant name corresponds to an existing customer; after determining that the applicant name corresponds to an existing customer, receiving, by the one or more processors, customer data from one or more sources indicating a current customer location;

determining, by the one or more processors, whether the current customer location corresponds to, or is within a predetermined distance of, a location of the source computer from which the virtual application originated; and

in response to determining that the current customer location does not correspond to the location of the source computer, generating, by the one or more processors, an electronic alert indicating such.

8. The computer-implemented method of claim 1, wherein determining an online browsing or search history associated with the IP address includes retrieving an online browsing or search history that is stored with or within the virtual application.

9. The computer-implemented method of claim 1, wherein determining an online browsing or search history associated with the IP address includes receiving, by the one or more processors and via the one or more transceivers, the online browsing or search history over one or more radio links.

10. The computer-implemented method of claim 1, wherein the virtual application is an application for financial services or products, a credit card application, or an application for auto, homeowners, or life insurance.

11. A computer system configured to use IP addresses and browsing activity to identify fraudulent virtual applications, the computer system comprising:

one or more transceivers; and

one or more processors configured to:

determine an applicant name on a virtual application; determine an IP address of a source computer from which the virtual application originated;

determine an online browsing or search history associated with the IP address;

determine whether the online browsing or search history indicates recent Internet searches for the applicant name; and

in response to determining that the online browsing or search history does indicate recent Internet searches for the applicant name, (i) flag the virtual application as fraudulent and (ii) generate an electronic alert indicating that the virtual application is fraudulent.

51

12. The computer system of claim 11, wherein the virtual application includes a field or routine that records recent online browsing or searching activity of the source computer, and wherein the received virtual application includes the recorded recent online browsing or searching activity of the source computer.

13. The computer system of claim 11, wherein the virtual application includes a field or routine that records recent social media searching activity of the source computer.

14. The computer system of claim 11, wherein the virtual application includes an IP address field of the source computer, and wherein the one or more processors are configured to determine the IP address of the source computer at least by reading the IP address field of the virtual application.

15. The computer system of claim 11, wherein the virtual application includes a field indicating a current GPS location of the source computer.

16. The computer system of claim 11, wherein the online browsing or search history associated with the IP address is stored with or within the virtual application.

17. A non-transitory, computer-readable medium storing instructions that, when executed by one or more processors, cause the one or more processors to:

determine an applicant name on a virtual application;
determine an IP address of a source computer from which the virtual application originated;

52

determine an online browsing or search history associated with the IP address; determine whether the online browsing or search history indicates recent Internet searches for the applicant name; and

in response to determining that the online browsing or search history does indicate recent Internet searches for the applicant name, (i) flag the virtual application as fraudulent and (ii) generate an electronic alert indicating that the virtual application is fraudulent.

18. The non-transitory, computer-readable medium of claim 17, wherein the virtual application includes a field or routine that records recent online browsing or searching activity of the source computer, and wherein the received virtual application includes the recorded recent online browsing or searching activity of the source computer.

19. The non-transitory, computer-readable medium of claim 17, wherein the virtual application includes a field or routine that records recent social media searching activity of the source computer.

20. The computer system of claim 17, wherein the virtual application includes an IP address field of the source computer, and wherein the instructions cause the one or more processors to determine the IP address of the source computer at least by reading the IP address field of the virtual application.

* * * * *